vmware by Broadcom

# VMware Aria Operations 8.18

# Table of Contents

# VMware Aria Operations 8.18.3 Release Notes

This document contains the following sections



## Introduction

VMware Aria Operations 8.18.3 | 30 January 2025

Check for additions and updates to these release notes.

## Build Details

VMware Aria Operations 8.18.3 | 2024 | Build 24521408

VMware Aria Operations Upgrade PAK | 2024 | Build 24521385

**Note:** *This upgrade .pak file includes the OS upgrade files from Photon to Photon, the vApp upgrade files, and Cloud Appliance upgrade files.*

VMware Aria Operations Upgrade PAK (Only Cloud Proxy) | 2024 | Build 24521385

**Note:** *This upgrade .pak file includes only the Cloud Appliance upgrade file.*

VMware Aria Operations 8.18.3 Pre-Upgrade Assessment Tool | 2024 | Build 24527447

VMware Aria Operations Cloud Appliance (VMware Aria Operations Cloud Proxy) 8.18.3 | 2024 | Build 24521398

Check frequently for additions and updates to these release notes.

**Document Revision History**

| Date | Edition | Changes |
|---|---|---|
| January 30, 2025 | 1 | Initial edition. |

## What's New

This release resolves CVE-2025-22222. For more information on this vulnerability and its impact on VMware products, please see VMSA-2025-0003. For details about defects that have been fixed, see KB 386227.

**Product Support Notice**

### Remote Collectors

Cloud proxies have been established as a long-term solution for data collection. In future releases, the introduction of new features will be exclusively limited to cloud proxies, if relevant. The announcement regarding the expiration of support for remote collectors and the prohibition of deploying new remote collectors has already been added in the 8.10 Release notes and accompanying documentation.

VMware Aria Operations 8.14 is the last release supporting remote collectors. In the 8.16 and later releases, upgrades will not be allowed if there are remote collectors in use. To upgrade to the next version or a later one, it is imperative to replace all remote collectors with cloud proxies.

Complete the following steps:

- If you are using a collector group, you must add new cloud proxies to it and remove the remote collectors.
- It is recommended that if you are not already using a collector group, create a new collector group and deploy cloud proxies into it. Reassign all adapter instances that were previously associated with remote collectors to the collector groups and remove the remote collectors.

### Deprecation of the XML Media Type from VMware Aria Operations REST APIs

The current REST APIs in VMware Aria Operations 8.18 support both JSON and XML types. In the next major release, new APIs or new functionalities of existing APIs will not support the XML type and will support only JSON. It is recommended that you use JSON for data interchange. However, support for the XML type in all existing APIs will continue.

### vRealize Application Remote Collector

vRealize Application Remote Collector is not supported from vRealize Operations 8.10 and above. Migrate all Telegraf end points to cloud proxy before upgrading to vRealize Operations 8.10 and above. For information about migrating from vRealize Application Remote Collector to cloud proxy, see KB 83059.

### VMware Cloud

- The VMware Cloud on Dell EMC adapter has been deprecated from VMware Aria Ops 8.18. As a result, the current deployment might not work.

### Native Public Cloud

The integrations for Amazon Web Services, Microsoft Azure, Oracle Cloud VMware Solution, and Google Cloud Platform are no longer available natively. They will be accessible via the Marketplace.

> **Note:** If you are upgrading from VMware Aria Operations 8.14.x to 8.18.3 and you have configured the Google Cloud Platform account, Google Cloud Platform will stop collecting data after you upgrade to VMware Aria Operations 8.18.3. To workaround this issue, the Google Cloud Platform adapter must be upgraded to version 8.18 immediately after cluster upgrade. To avoid data loss, before you start cluster upgrade, ensure that the Google Cloud Platform 8.18 management pack is available in the Marketplace.

### What's New in This Release

For what's new content in 8.18.3 see the VMware Aria Operations 8.18 release notes.

### Metrics and Properties Modifications

The following KB article describes all the metrics and properties that have been modified in VMware Aria Operations 8.18:

Metrics added in VMware Aria Operations 8.18

### Instanced Metrics

Instanced metrics are deactivated by default after deploying or upgrading to vRealize Operations 8.2 or later, and after importing a policy from older versions. To re-activate instanced metrics in vRealize Operations 8.2 or later, see KB 81119.

### Basic Authentication

Basic authentication using the REST API is deprecated and deactivated in VMware Aria Operations 8.18 fresh deployments by default. Instances that have been upgraded to VMware Aria Operations 8.18.3, will inherit the same properties before the upgrade. It is recommended that you use token-based authentication instead. If you still need to activate or deactivate basic authentication, see KB 77271.

### Active Directory Authentication Sources

Logging in to VMware Aria Operations with a short name will be successful only if the user name's domain suffix matches the domain name specified in the **Base DN** option. Otherwise, the full user name with the domain suffix is required during login. For more information, see KB 68131.

## System Requirements

Review this section before you install or update VMware Aria Operations.

Sizing and Scaling

The CPU, memory, and disk requirements that meet the needs of a particular environment depend on the number and type of objects in your environment and data collected. This includes the number and type of adapters installed, the use of HA (High Availability) and CA (Continuous Availability), the duration of data retention, and the quantity of specific data points of interest. VMware updates Knowledge Base article 2093783 with the most current information about sizing and scaling. The Knowledge Base article includes overall maximums and spreadsheet calculations that provide a recommendation based on the number of objects and metrics you expect to monitor.

Deployment Formats

You can deploy VMware Aria Operations 8.18.3 with VMware virtual appliance.

If you are deploying a VMware Aria Operations virtual appliance and VMware Aria Operations Cloud Appliance (cloud proxy), use a VMware vSphere Client to connect to a VMware vCenter Server, and deploy the virtual appliance through the vCenter Server instance. The VMware Aria Operations virtual appliance and VMware Aria Operations Cloud Appliance (cloud proxy) must be deployed on hosts that are:

- ESX/ESXi 6.5 Update 1 or later and managed by VMware vCenter Server 6.5 or later.
- If you have VMware Aria Operations virtual appliance deployed on ESXi 6.0 or older hosts, you must first upgrade vCenter Server and ESXi to version 6.5 Update 1 or later, and then upgrade to VMware Aria Operations 8.18.3.

## Hardware Versions, Cipher Suites and Protocols, and Log4j

### Hardware Versions

The minimum hardware version required for VMware Aria Operations 8.x releases is version 11. If your VMware Aria Operations virtual appliance had a hardware version earlier than 11, you must first upgrade to hardware version 11 on VMware Aria Operations virtual appliance and then upgrade to VMware Aria Operations 8.18.3.

### Cipher Suites and Protocols

For information about cipher suite lists and relevant protocols, see Cipher Suites and Protocols.

**Log4j version** is at 2.17.2.

## VMware Product Compatibility

VMware Product Compatibility

**Note:** The VMware Product Interoperability Matrix provides details about the compatibility of VMware Aria Operations with VMware products.

**Note:** For FIPS mode compatibility details, see the footnotes in the interoperability matrix. The product will not work in FIPS mode if there are footnotes that state that it will not work in FIPS mode.

## Solutions and Browser Support

Solutions Support

In addition to the VMware solutions (vSphere, VMware Aria Operations for Logs, vSAN, Service Discovery, NSX-T, and many more), see the VMware Marketplace for many more solutions. These solutions work with Virtual Appliance single or multiple nodes.

Browser Support

This VMware Aria Operations release supports all current Web browsers, although only the following browsers were tested with this release:

- Google Chrome: Versions 132 and 131
- Mozilla Firefox: Versions 134.0.2 and 133.0.3
- Microsoft Edge: Versions 132 and 131
- Safari: 17.3 and 17.4

**Note:** Support for Internet Explorer has been dropped from vRealize Operations 8.4 onwards.

## SDDC Compliance

Ensure compliance of your vSphere, vSAN, and NSX-T resources deployed across private data centers, and other VMware managed clouds such as VMware Cloud Foundation (VCF), VMware Cloud on AWS, Azure VMware Solution, Google Cloud VMware Engine, Oracle Cloud VMware Solution, and VMware Cloud on Dell EMC using various compliance packs in VMware Aria Operations. For information about supported versions of vSphere, vSAN, and NSX, see VMware SDDC Benchmark Details.

## Installing and Upgrading VMware Aria Operations

**Upgrading to VMware Aria Operations 8.18.3 or upgrading any Management Pack (MP), resets out-of-the-box content as part of the software upgrade process. This implies that the user modifications made to default content such as alert definitions, symptom definitions, recommendations, policies, views, dashboards, widgets, and reports are overwritten. You need to clone or backup the content before you upgrade to VMware Aria Operations 8.18.3.**

**Notes:**

- To see the supported direct upgrade path, refer to the Product Interoperability Matrix.
- While upgrading to VMware Aria Operations 8.18.3, the expected size of the /dev/sda for Photon OS is 20 GB (hard disk 1). For information about this requirement, see KB 75298.
- It is always recommended to run the Pre-Upgrade Assessment tool before an upgrade. A pre-upgrade assessment report that is generated will provide you with the recommended replacements. This tool provides you with an impact analysis following the reduction of metrics in various versions of the product. For more details on using the Pre-Upgrade Assessment Tool, see KB 369264.

- It is mandatory to create a snapshot of each node in a cluster before you update a VMware Aria Operations cluster. Once the update is complete, you must delete the snapshot to avoid performance degradation.

VMware Aria Suite Lifecycle 8.18 supports the installation of VMware Aria Operations 8.18. For more information, see the VMware Aria Suite Lifecycle 8.18 Release Notes.

Refer to the VMware Aria Operations Upgrade Center that has information about upgrading VMware Aria Operations. Refer to the Product Matrix for information about supported versions of VMware Aria Operations.

The VMware Aria Operations Information Center has detailed information about installation and software updates.

Refer to the Getting Started with VMware Aria Operations Guide that provides guidance for a VMware Aria Operations virtual appliance, before you install or update VMware Aria Operations.

Deploy vSphere with Operations Management (any edition) and VMware Aria Operations Standard together in one deployment.

Deploy vCloud Suite/VMware Aria Suite Standard, Advanced, or Enterprise and VMware Aria Operations Advanced or Enterprise edition together in one deployment.

**Note:** You can also install VMware Aria Operations by using VMware Aria Suite Lifecycle Manager.

## Resolved Issues

For details on resolved issues, see KB 386227.

## Known Issues

## Installation and Upgrade Issues

**Upgrade to VMware Aria Operations 8.18.3 will impact some management packs**
When you upgrade to VMware Aria Operations 8.18.3, some management packs might not be compatible with JDK11 and would require an upgrade to a JDK11 compatible version. Contact the vendor to confirm if the management pack is compatible with VMware Aria Operations 8.18.3.
**Workaround:** See KB 89675 for more details.

**Upgrade might fail if any of the nodes are running on VMs with US/Pacific-New timezone**
PostgreSQL database systems no longer support the US/Pacific-New timezone, which was just an alias for America/Los_Angeles timezone. If any of the VMware Aria Operations nodes are running on VMs with US/Pacific-New timezone, upgrade might fail.
**Workaround:** Change VM timezones from US/Pacific-New to America/Los_Angeles, and then upgrade.

**After upgrading vrops org (8.4 to 8.5), content upgrade is failing in Cloud and On-prem.**
**After you upgrade vRealize Operations from 8.4 to a later release, content upgrade and agent management actions fail**

After you upgrade from a previous release, content upgrade and agent management actions fail on on-prem and SaaS when you also upgrade cloud proxy.
**Workaround:**

Complete the following steps:

1. SSH to the Cloud Proxy VM.

2. Run the following command: */rpm-content/ucp/subsequentboot.sh*

You can view the log from the following location: */opt/vmware/var/log/ucp-subsequentboot*

## General Issues

**Agent Install Fails for Non-HA Collector Groups**
Product managed agent install fails with an error for non-HA collector groups

**Workaround:**

- Move the cloud proxy out of the non-HA collector group so that it functions as a standalone cloud proxy, or
- Activate the collector group so that it functions as an HA collector group.

**Installation of agents in bulk on product managed VMs, fail with an error**
Installing agents in bulk on product managed VMs may fail for some VMs with the following error: "Exception occurred while trying to upload the command".
**Workaround:** Reattempt the installation of the agent on those VMs where installation failed. This often resolves the problem and the agents are installed successfully.

**The Kubernetes management pack does not work after an upgrade to VMware Aria Operations 8.18.3**
If you have the Kubernetes management pack installed and have upgraded to VMware Aria Operations 8.18.3, the management pack does not work.
**Workaround:** It is recommended that you upgrade the management pack to version 2.2.

**The Google Cloud Platform cloud management pack fails after you upgrade from VMware Aria Operations 8.14.x to VMware Aria Operations 8.18.3**
If you have configured the Google Cloud Platform management pack on VMware Aria Operations 8.14.x and have upgraded to VMware Aria Operations 8.18.3, the management pack will not work.

**The CMA management pack has broken links in many of its dashboard widgets**
The CMA management pack is no longer updated and will be removed in future releases. It is currently available, but not fully functional, and in a deactivated state. You can activate the management pack if required, but most functionality is incomplete.

**Workaround:** None

**VMware Chargeback migration must be run with an "admin" user**
While migrating from VMware Chargeback to VMware Aria Operations, you must run the migration with an "admin" user.
**Workaround:** None

**VC Pricing Policies migration must be run with an "admin" user**
While migrating from VC Pricing card to VMware Aria Operations policies, you must run the migration with an "admin" user.
**Workaround:** None

**Side-details panel for the collector group fails to load if many cloud proxies are configured**
**The side-details panel for the collector group fails to load if more than 90 cloud proxies are configured.**
**Workaround:** Configure less than 90 cloud proxies.

**vSAN Oversubscription Capacity and Oversubscription Ratio metrics are incorrect**

Oversubscription Capacity and Oversubscription Ratio metrics calculation is incorrect for the vSAN ESA environment. The Oversubscription metrics are applicable only for the vSAN OSA clusters.
**Workaround:** None

**Agent install fails with the error: Connect to salt master**
When you install an agent to monitor applications using Telegraf, the following error may occur: *`Connect to Salt Master`*. The error occurs because the end point VM is unable to connect with the salt master.
**Workaround:** Verify the ownership of the folders in */ucp/salt/pki/master* in cloud proxy. The ownership must be admin/admin. If the ownership is not admin, reset permissions/ownership.

1. Run the `/ucp/ucp-config-scripts/ucp-firstboot.sh` script and verify the permissions of the folder.
2. Reinstall the agent after successful execution of the script.

Agent install should be successful.

**VMware Aria Operations API (Suite-API) based Telegraf agent installation fails with an error**
When you use an API end point for Telegraf agent installation in VMware Aria Operations API (Suite-API) using `/api/applications/agents`, the following error occurs:

```
VM with ID - id is not connected to any ARC or Cloud Proxy
```

```
Ex- VM with ID - 94b8e2eb-37fa-4a44-b241-b50e3a013bf7 is not connected to any ARC or Cloud Proxy
```

**Workaround:**

1. Move the vCenter adapter collector/group back to the same cloud proxy/collector group that was selected the during the first Telegraf agent installation of VMs for that vCenter.
2. Navigate to **Data Sources > Integrations > Accounts > vCenter** and select the vCenter where the endpoint is located.
3. Click **Edit** and under the Cloud Account Information, select the **vCenter** tab **> Collector/Group** and change collector/collector group from drop down options and then click **Save**.
4. After the Telegraf agent installation is successful, move the vCenter adapter collector/group back to the initial cloud proxy/collector group.

**JBoss server running in domain mode does not support LCM using custom Telegraf**
You cannot deploy jolokia.war across all the servers that run as a part of the JBoss domain mode and hence metrics cannot be collected.
**Workaround:** Monitoring domain mode configuration in JBoss is supported in open source Telegraf. You can deploy jolokia.war across multiple servers.

**The Add button in the Manage Telegraf Agents page remains activated after the creation of an application service**
While monitoring specific application services using Telegraf, the Add button in the Manage Telegraf Agents page is active after the creation of the second instance of the following application services: MSIIS, Active Directory, SharePoint, MSExchange, and Network Time Protocol.
**Workaround:** Delete the previous configurations of the specific application services and configure a new one.

**Concurrent activation or deactivation of plugins during application monitoring using suite-api does not work**
During application monitoring using suite-api, when you activate or deactivate plugins concurrently, plugin activation or deactivation does not work. An exception is logged in the ucpapi.log file.
**Workaround:** Provide a gap of one second between each thread during activation or deactivation of the plugins using suite-api.

**When you get an application instance's configuration status using a Rest-API call, a wrong status is returned when the same configuration was installed or uninstalled previously**

If you installed or uninstalled an application instance and then uninstalled or installed it respectively, the API returns "SUBMITTING" when you try to get the application instance install or uninstall configuration status using "GET /api/applications/agents/services/{taskId}/status".

**Workaround:** While getting the application instance configuration status using the "GET /api/applications/agents/services/{taskId}/status" API, ensure that you use the latest performed task ID. The result of calls with older IDs is undefined.

**When you upgrade from vRealize Operations 8.1 to VMware Aria Operations 8.18.3, the Cassandra application service is displayed as Java Application in the Services Discovered/Configured column of the Manage Telegraf Agents tab**

During application monitoring, after you upgrade from vRealize Operations 8.1 to VMware Aria Operations 8.18.3, the Cassandra application service is displayed as Java Application in the Services Discovered/Configured column of the Manage Telegraf Agents tab. This happens if the Cassandra application service is configured for monitoring in vRealize Operations 8.1, using the vRealize Application Remote Collector generic Java plugin.

**Workaround:** Deactivate the Java plugin before you upgrade the vRealize Application Remote Collector agent on the VM. After upgrading the vRealize Application Remote Collector agent, the Cassandra plugin will be discovered and can then be activated.

**System language settings affect service discovery**

Service discovery might not work if the system language is different from English. For languages different from English, network connection state values might be different from constants defined in standard RFCs.

**Workaround:** None

**Reconfigure the Project Price widget if the VMware Aria Automation integration is deactivated and then reactivated**

There is data missing in the **Project Price** widget of the **Cloud Automation Project Price Overview** dashboard when you deactivate and then reactivate an existing VMware Aria Automation integration.

**Workaround:** Reconfigure the **Project Price** widget by completing the following steps:

1. After you deactivate and then reactivate the VMware Aria Automation integration, navigate to **Dashboards** > **VMware Aria Automation** > **Cloud Automation Project Price Overview** dashboard.
2. Edit the **Project Price** widget.
3. Navigate to the **Input data** section.
4. Select the **'+' (Add Object)** button and select the CAS World object from **VMware Aria Automation > CAS World**.
5. Click **OK**.
6. Navigate to the **Output data** in the same widget configuration mode. Search for and select the CAS Project Price View object from the list.
7. Save the widget.

**While monitoring applications, you cannot activate a plugin with the same fields till the plugin configuration is deleted**

An error message is displayed in the user interface of VMware Aria Operations that states the following: 'Failed to update resource: Resource with same key already exists'.

**Workaround:** Manually delete the existing plugin configuration and then continue with the activation of the plugin. If the problem persists, delete the corresponding resource from the inventory.

**Alerts from the vSAN adapter that correspond to vSAN health check tests are not canceled if the test is removed from the vSAN Health Service.**

VMware Aria Operations cannot detect and cancel deleted alerts.

**Workaround**: Manually cancel the alert from the user interface of VMware Aria Operations.

**Despite deleting ucp-adapter instance certificates, users can run actions**
If users delete ucp-adapter instance certificates, they can still run actions such as, start and stop an agent, configure remote checks, and so on.
**Workaround:** None

**The compliance score for a user with limited object visibility is the same as for a user with complete object visibility**
The compliance score is calculated for objects that are not visible (not assigned) to the current user.
**Workaround:** Complete the following steps:

1. Create a custom group with objects visible (assigned) to the user.
2. For that group, apply a policy in which the needed set of compliance alert definitions is activated.

   If that set is activated only in one active policy (the one that is applied to the custom group), the compliance benchmark based on those alert definitions will display the correct score.

## User Interface Issues

**After upgrading cloud proxy and VMware Aria Operations to 8.18.3, the VM name is not appended to the MSSQL instance**
After upgrading cloud proxy and VMware Aria Operations to 8.18.3, the VM name is not appended to the MSSQL instance for existing or newly activated plugins.
**Workaround**: None

**The Last Year option in the date picker is not intuitive**
The Last Year option in the date picker indicates that the time range starts from the end of the previous month and goes back a year. It does not indicate a time range that spans one year from the current date or the whole previous year.
**Workaround**: None

# Known Issues

## Installation and Upgrade Issues

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

## General Issues

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

### VMware Aria Operations 8.18.2 Release Notes

## VMware Aria Operations 8.18.2 Release Notes

## VMware Aria Operations 8.18.2 Release Notes

## VMware Aria Operations 8.18.2 Release Notes

## User Interface Issues

## VMware Aria Operations 8.18.2 Release Notes

## VMware Aria Operations 8.18.2 Release Notes

# Resolved Issues

## VMware Aria Operations 8.18.1 Release Notes

## VMware Aria Operations 8.18.1 Release Notes

# Known Issues

## Installation and Upgrade Issues

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

## General Issues

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

### VMware Aria Operations 8.18.1 Release Notes

## VMware Aria Operations 8.18.1 Release Notes

## VMware Aria Operations 8.18.1 Release Notes

## VMware Aria Operations 8.18.1 Release Notes

## User Interface Issues

## VMware Aria Operations 8.18.1 Release Notes

## VMware Aria Operations 8.18.1 Release Notes

# Resolved Issues

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

# Known Issues

### Installation and Upgrade Issues

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### General Issues

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

### VMware Aria Operations 8.18 Release Notes

## VMware Aria Operations 8.18 Release Notes

## VMware Aria Operations 8.18 Release Notes

## VMware Aria Operations 8.18 Release Notes

## VMware Aria Operations 8.18 Release Notes

## VMware Aria Operations 8.18 Release Notes

## User Interface Issues

## VMware Aria Operations 8.18 Release Notes

## VMware Aria Operations 8.18 Release Notes

# VMware Aria Operations Reference Architecture (8.18)

The *VMware Aria Operations Architecture Guide* provides recommendations for deployment topology, hardware requirements, interoperability, and scalability for VMware Aria Operations.

For information about software requirements, installation, and supported platforms see the VMware Aria Operations documentation.

## Best Practices for Deploying VMware Aria Operations

Implement all the best practices when you deploy a production instance of VMware Aria Operations.

### Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

> **NOTE**
> The master node is now referred to as the primary node. The master replica node is now referred to as the primary replica node.

- Deploy analytics nodes in the same vSphere Cluster except when activating Continuous Availability.
- Deploy analytics nodes with the same disk size on storage of the same type.
- When activating Continuous Availability, separate analytics nodes into fault domains based on their physical location.
- Depending on the size and performance requirements for analytics nodes, apply Storage DRS Anti-Affinity rules to ensure that nodes are on separate datastores.
- Set Storage DRS to manual for all VMware Aria Operations analytics nodes.
- If you deploy analytics nodes into a highly consolidated vSphere cluster, configure the resource reservation to ensure optimal performance. Ensure that the virtual CPU to physical CPU ratio is not negatively impacting the performance of analytics nodes by validating CPU ready time and CPU co-stop.
- Analytics nodes have a high number of vCPUs to ensure performance of the analytics computation that occurs on each node. Monitor CPU Ready time and CPU Co-Stop to ensure that analytics nodes are not competing for CPU capacity.
- If the sizing guideline provides several configurations for the same number of objects, use the configuration which has the least number of nodes. For example, if the number of collecting is 120,000, configure the cluster with four extra-large nodes instead of 12 large nodes.
- Deploy an extra even number of nodes to activate Continuous Availability. If the current configuration is an odd number of analytics nodes, deploy an extra analytics node to create an even pairing.

### Witness Nodes

A witness node is required when continuous availability is activated to manage the analytics nodes in the fault domains. VMware Aria Operations can have only one witness node in its cluster.

- Deploy the witness node before activating continuous availability.
- Deploy the witness node using the witness configuration.
- Deploy the witness node in a different cluster separate from the analytics nodes.

### Cloud Proxy

Using cloud proxies in VMware Aria Operations, you can collect and monitor data from your remote data centers. You can deploy one or more cloud proxies in VMware Aria Operations to create a one-way communication between your remote environment and VMware Aria Operations. The cloud proxies work as one-way remote collectors and upload data from the remote environment to VMware Aria Operations. Cloud proxies can support multiple vCenter Server accounts.

## Cloud Proxy and Telegraf Agents

- Deploy Cloud Proxy in the same vCenter as the end point VMs on which you want to deploy the Telegraf agents. For Cloud Proxy sizing information see the Sizing Guidelines (KB 2093783).
- Ensure that your operating system platform is supported by Cloud Proxy, and the most recent versions of Windows and Linux OS are supported.
- System times must be synchronized between cloud proxy, end point VMs, the vCenter, ESX host, and VMware Aria Operations. To ensure synchronized time, use Network Time Protocol (NTP).
- Ensure that all the prerequisites are met. For more information, see Prerequisites.
- Disable UAC on Endpoint VMs before installing the Telegraf agent. If you cannot do this due to security restrictions, see KB 70780 for a work around script.
- Ensure that the version later than 12.3.5 of VMware Tools is installed on the end point VM on which you want to deploy the Telegraf agent.
- To deploy Telegraf agents onto end point VMs, ensure that the following prerequisites are met for the user account being used for deployment:
  Windows - The user account must be either:
  - An administrator account
  - A non-administrator account that is a member of the built-in administrator group

  Linux - The user account must be either:

  – A root user with all privileges
  – A non-root user with all privileges
  – A non-root user with specific privileges

  For more information, see User Account Prerequisites in the *VMware Aria Operations Configuration Guide*.

## Management Packs and Adapters

Various management packs and adapters have specific configuration requirements. Ensure that you are familiar with all prerequisites before you install a solution and configure the adapter instance.

- Utilize collector groups to separate data collection into fault domains when continuous availability is activated.

## Deployment Formats

Deploy VMware Aria Operations with the same VMware Aria Operations version for the following node types:

- Primary
- Primary Replica
- Data
- Witness

See the topic, *Installing VMware Aria Operations* in the *Getting Started with VMware Aria Operations* guide for more information.

## Initial Considerations for Deploying VMware Aria Operations

For the production instance of VMware Aria Operations to function optimally, your environment must conform to certain configurations. Review and familiarize yourself with these configurations before you deploy a production instance of VMware Aria Operations.

### Sizing

For information about the number of monitored resources and how many analytics nodes are supported by VMware Aria Operations, refer to the Sizing Guidelines (KB 2093783). You must plan sizing of your VMware Aria Operations instance to ensure performance and support.

### Environment

Deploy analytics nodes in the same vSphere cluster and use identical or similar hosts and storage. If you cannot deploy analytics nodes in the same vSphere cluster, you must deploy them in the same geographical location.

When continuous availability is activated, deploy analytics nodes in fault domains in the same vSphere cluster and use identical or similar hosts and storage. Fault domains are supported on vSphere stretched clusters.

Analytics nodes must be able to communicate with one another always. The following vSphere events might disrupt connectivity.

- vMotion
- Storage vMotion
- High Availability (HA)
- Distributed Resource Scheduler (DRS)

Due to a high level of traffic between analytics nodes, all analytics nodes must be on the same VLAN and IP subnet, and that VLAN is not stretched between data centers, when continuous availability is not activated.

When the continuous availability is activated analytics nodes within each fault domain should be on the same VLAN and IP subnet, and they should be able to communicate with each other.

Latency between analytics nodes cannot exceed 5 milliseconds, except when continuous availability is activated, where latency between fault domains cannot exceed 10 milliseconds but analytics nodes, within each fault domain, still cannot exceed 5 milliseconds. The bandwidth must be equal to or faster than 10 GB per second.

If you deploy analytics nodes into a highly consolidated vSphere cluster, configure resource reservations. See the Sizing Guidelines (KB 2093783) for more information. If you experience performance issues, review the CPU ready and co-stop to determine if the virtual to physical CPU ratio is the cause of the issues. For more information about how to troubleshoot VM performance and interpret CPU performance metrics, see KB 1017926.

You can deploy the witness node behind a firewall. You cannot use NAT between the witness node and analytics nodes.

### Multiple Data Centers

VMware Aria Operations can be stretched across data centers only when continuous availability is activated. The fault domains may reside in separate vSphere clusters; however, all analytics nodes across both fault domains must reside in the same geographical location.

For example, the first data center is located in Palo Alto but is configured in two different buildings or in different locations of the city (downtown and mid-town) will have latency that is less than 5 milliseconds. The second data center is located in Santa Clara so the latency between the two data centers is greater than 5 milliseconds but less than 10 milliseconds. See the Sizing Guidelines (KB 2093783) for network requirements.

If VMware Aria Operations is monitoring resources in additional data centers, you must use cloud proxies and deploy the cloud proxies in the remote data centers. You might need to modify the intervals at which the configured adapters on the cloud proxies collect information depending on latency.

It is recommended that you monitor collections to validate that they are completing in less than five minutes. See the Sizing Guidelines (KB 2093783) for latency, bandwidth and sizing requirements. If all requirements are met and collections are still not completing within the default 5 minutes time limit, increase the interval to 10 minutes.

### Certificates

A valid certificate signed by a trusted Certificate Authority, private, or public, is an important component when you configure a production instance of VMware Aria Operations. Configure a Certificate Authority signed certificate against the system before you configure agents.

You must include all analytics nodes, witness nodes, and load balancer DNS names in the Subject Alternative Names field of the certificate.

### Adapters

It is recommended that you configure adapters to cloud proxies in the same data center as the analytics cluster for large and extra-large deployment profiles. Configuring adapters to cloud proxies improves performance by reducing load on the analytics node. As an example, you might decide to configure an adapter to cloud proxies if the total

resources on a given analytics node begin to degrade the node's performance. You might configure the adapter to a large cloud proxy with the appropriate capacity.

Configure adapters to cloud proxies when the number of resources the adapters are monitoring exceeds the capacity of the associated analytics node.

**Authentication**

You can use the Platform Services Controller for user authentication in VMware Aria Operations. For more information about deploying a highly available Platform Services Controller instance, see Deploying the vCenter Server Appliance in the *VMware vSphere Documentation*. All Platform Services Controller services are consolidated into vCenter, and deployment and administration are simplified.

**Load Balancer**

For more information about load balancer configuration, see the VMware Aria Operations Load Balancing guide.

## Scalability Considerations

Configure your initial deployment of VMware Aria Operations based on the anticipated use.

For more information about sizing, see the Sizing Guidelines (KB:2093783).

**Analytics Nodes**

Analytics nodes consist of a primary node, a primary replica node, and data nodes.
For enterprise deployments of VMware Aria Operations, deploy all nodes as medium, large or extra-large deployments, depending on sizing requirements and your available resources.

**Scaling Vertically by Adding Resources**

If you deploy analytics nodes in a configuration other than large, you can reconfigure the vCPU and memory. It is recommended to scale up the analytics nodes in the cluster before scaling out the cluster with additional nodes. VMware Aria Operations supports various node sizes.

**Scaling Vertically by Increasing Storage**

You can increase storage independently of vCPU and Memory.

To maintain a supported configuration, data nodes deployed in the cluster must be the same node size.

For more information about increasing storage, see the topic, Add Data Disk Space to a Aria Operations vApp Node in the *Getting Started* guide. You cannot modify the disks of virtual machines that have a snapshot. You must remove all snapshots before you increase the disk size.

**Scaling Horizontally (Adding nodes)**

To see the number of t extra-large analytics nodes in a cluster, or the number of extra-large nodes in a cluster when continuous availability is activated, see the Sizing Guidelines (KB:2093783).

To maintain a supported configuration, analytics nodes deployed in the cluster must be the same node size.

**Witness Node**

VMware Aria Operations provides a single size regardless of the cluster size since the witness node does not collect nor process data.

**Remote Collectors**

> **NOTE**
> Fresh deployment of remote collectors is not supported in VMware Aria Operations, starting from version 8.10. Remote collectors are available only if you had deployed them in a previous version of VMware Aria Operations. If you require a new agent to collect data, you must deploy a cloud proxy. For more information on how to deploy a cloud proxy, see the topic, Installing Cloud Proxy in the Getting Started guide.

VMware Aria Operations supports two sizes for remote collectors, standard and large. The maximum number of resources is based on the aggregate resources that are collected for all adapters on the remote collector. In large

scale VMware Aria Operations monitored environment, you might experience a slow responding UI, and metrics are slow to be displayed.

**Cloud Proxy**

VMware Aria Operations supports two sizes for Cloud Proxy, small and large. The maximum number of resources is based on the aggregate resources that are collected for all adapters on the Cloud Proxy. In large scale VMware Aria Operations monitored environment, you might experience a slow responding UI, and metrics are slow to be displayed. Install a remote collector Cloud Proxy in areas when the latency is more than what is prescribed in the Sizing Guidelines. See (KB 2093783) for more information.

## High Availability Considerations

High availability creates a replica for the VMware Aria Operations primary node and protects the analytics cluster against the loss of a node.

**Cluster Management**

Clusters consist of a primary node, a primary replica node, and data nodes.

Activating High Availability within VMware Aria Operations is not a disaster recovery solution. When you activate High Availability, information is stored (duplicated) in two different analytics nodes within the cluster. This doubles the system's compute and capacity requirements. If either the primary node or the primary replica node is permanently lost, then you must deactivate, and then reactivate High Availability to reassign the primary replica role to an existing node. This process, which includes a hidden cluster rebalance, can take a long time.

**Analytics Nodes**

Analytics nodes consist of a primary node, primary replica node, and data nodes.

When you activate High Availability, you protect VMware Aria Operations from data loss when only a single node is lost. If two or more nodes are lost, there may be permanent data loss. Deploy each analytics node to separate hosts to reduce the chance of data loss if a host fails. You can use DRS anti-affinity rules to ensure that the VMware Aria Operations nodes remain on separate hosts.

**Collector Group**

In VMware Aria Operations, you can create a collector group. A collector group is a collection of nodes (cloud proxy, and analytics nodes). You can assign adapters to a collector group, rather than assigning an adapter to a single node.

**NOTE**
A collector group must contain the same type of nodes. You cannot mix cloud proxy, and analytics nodes in a collector group.

If the node running the adapter fails, the adapter is automatically moved to another node in the collector group.

Assign all normal adapters to collector groups, and not to individual nodes. Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

For more information about adapters, see Adapter and Management Packs Considerations.

## Continuous Availability Considerations

Continuous Availability (CA) separates the VMware Aria Operations cluster into two fault domains and protects the analytics cluster against the loss of a fault domain.

### Cluster Management

Clusters consist of a primary node, a primary replica node, a witness node, and data nodes.

Activating Continuous Availability within VMware Aria Operations is not a disaster recovery solution.

When you activate Continuous Availability, information is stored (duplicated) in two different analytics nodes within the cluster but stretched across fault domains. Due to sizing requirements, continuous availability requires doubling the system's compute and capacity requirements.

If either the primary node or primary replica node is permanently lost, then you must replace the lost node, which will become the new primary replica node. If it is necessary to have the new primary replica node as the primary node, then you can take the current primary node offline and wait until the primary replica node is promoted to the new primary node. Then bring the former primary node back online and it will be the new primary replica node.

## Fault Domains

Fault domains consist of analytics nodes, separated into two zones.

A fault domain consists of one or more analytics nodes grouped according to their physical location in the data center. When configured, two fault domains allow VMware Aria Operations to tolerate failures of an entire physical location and failures from resources dedicated to a single fault domain.

## Witness Node

Witness node is a member of the cluster but not part of the analytics nodes.

To activate CA within VMware Aria Operations, deploy the witness node in the cluster. The witness node does not collect nor store data.

The witness node serves as a tiebreaker when a decision must be made regarding availability of VMware Aria Operations when the network connection between the two fault domains is lost.

## Analytics Nodes

Analytics nodes consist of a primary node, primary replica node, and data nodes.

When you activate continuous availability, you protect VMware Aria Operations from data loss if an entire fault domain is lost. If node pairs are lost across fault domains, there may be permanent data loss.

Deploy analytics nodes, within each fault domain, to separate hosts to reduce the chance of data loss if a host fails. You can use DRS anti-affinity rules to ensure that the VMware Aria Operations nodes remain on separate hosts.

## Collector Group

In VMware Aria Operations, you can create a collector group. A collector group is a collection of nodes (Cloud Proxy, and analytics nodes). You can assign adapters to a collector group, rather than assigning an adapter to a single node.

> **NOTE**
> A collector group must contain the same type of nodes. You cannot mix Cloud Proxy, and analytics nodes in a collector group.

When activating continuous availability, collector groups can be created to collect data from adapters within each fault domain.

Collector groups do not have any correlation with fault domains. The functionality of a collector group is to collect data and provide it to the analytics nodes, which then VMware Aria Operations decides how to keep the data.

If the node running the adapter collection fails, the adapter is automatically moved to another node in the collector group.

Theoretically, you can install collectors in any place, provided the networking requirements are being met. However, from a failover perspective, it is not recommended to put all the collectors within a single fault domain. If all the collectors are directed to a single fault domain, VMware Aria Operations stops receiving data if a network outage occurs affecting that fault domain.

Assign all normal adapters to collector groups, and not to individual nodes. Hybrid adapters require a two-way communication between the adapter and the monitored endpoint.

For more information about adapters, see Adapter and Management Packs Considerations.

## Continuous Availability FAQs

With the introduction of continuous availability in VMware Aria Operations 8, there have been several frequently asked questions. This section is to help increase awareness and knowledge about continuous availability.

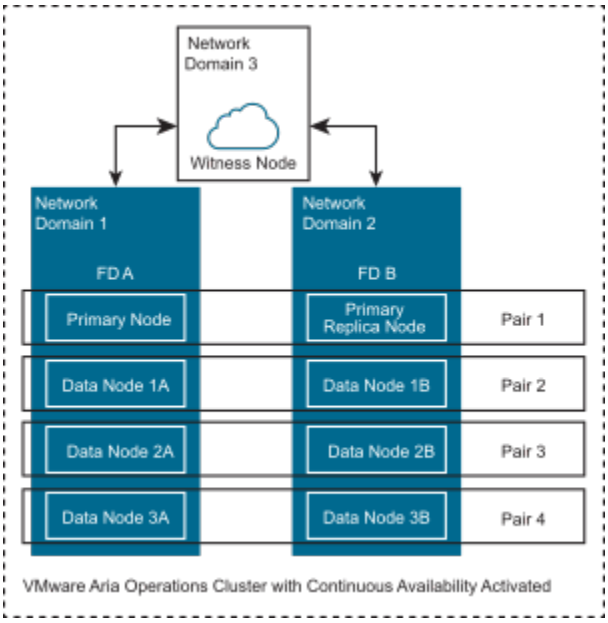


**How is the data stored in analytics nodes?**

When an object is discovered, VMware Aria Operations determines which node to keep the data, then copies (duplicates) the data to its pair node in the other fault domain. Every object is stored in two analytics nodes (node pairs) across the fault domains and they are always synchronized.

As an example, VMware Aria Operations has eight analytics nodes, CA is activated, and as a result each fault domain has four analytics nodes (see above diagram).

When a new object is discovered, VMware Aria Operations decides to store the data in “Data Node 2B” (primary) and automatically a copy of the data will be saved in “Data Node 2A” (secondary).

If somehow “FD A” becomes unavailable, then “primary” data from “Data Node 2B” will be used.

If somehow “FD B” becomes unavailable, then “secondary” data from “Data Node 2A” will be used.

**Which situations break a continuous availability cluster? Simultaneously losing the primary node or primary replica node and data nodes, or two or more data nodes in both fault domains, are not supported.**

Each analytics node from fault domain 1 has its node pair in fault domain 2 or vice versa.

Using the previously mentioned example, we will have four node pairs:

Primary + Replica Node

Data Node 1A (FD A) + Data Node 1B (FD B)

Data Node 2A (FD A) + Data Node 2B (FD B)

Data Node 3A (FD A) + Data Node 3B (FD B)

The two nodes of each node pair are always synchronized and storing the same data. Hence, the cluster continues to function without data loss while one node from all node pairs is available.

**What happens if one data node from one of the fault domains becomes unavailable?**

The cluster will be in a degraded state but continue to operate when one node becomes unavailable in either fault domain. There will be no data loss. The data node must be repaired or replaced so the cluster does not remain in a degraded state.

**Will the cluster break if two data nodes in fault domain 1 and the primary replica node in fault domain 2 are lost?**

In this example, the cluster will continue to work without data loss. If one analytics node from each node pair is still available, there will be no data loss.

**What happens if an entire fault domain becomes unavailable?**

The cluster will be in a degraded state but continue to operate when an entire fault domain becomes unavailable. There will be no data loss. The fault domain must be repaired and brought online so the cluster does not remain in a degraded state.

It the fault domain is unrecoverable, it is possible to replace the entire fault domain with newly deployed nodes. From the admin UI, only the primary replica node can be replaced. If the entire fault domain for the primary node is lost, you will need to wait until the primary node failover occurs and the primary replica node has been promoted as the new primary node.

**What is the proper process to re-add a failed node to a fault domain? How long will it take to sync up?**

The recommended procedure to re-add a failed node is to use the "Replace nodes of cluster" functionality within the admin UI. Once the replacement node has been added, the data will be synced. The sync time strongly depends on the object count, historical period of the objects, network bandwidth, and the load on the cluster.

**What happens when network latency between fault domains exceeds 20 ms? How long can VMware Aria Operations tolerate extended latency?**

Adhering to latency requirements is necessary to achieve optimal performance. The latency between fault domains should be < 10 ms, with peaks up to 20 ms during 20 sec intervals. For more information about network latency guidelines, see the Sizing Guidelines (KB 2093783).

**When network latency between fault domains goes above “20 ms during 20 sec intervals” for some period, but then recovers back to under 10 ms, how long does it take to resynchronize?**

High latency does not mean that synchronization has stopped. When an object is discovered, VMware Aria Operations will decide which node needs to keep the data (primary), then a second copy of the data will go to its node pair (secondary). Every object is stored in two analytics nodes (pairs) across both fault domains. Synchronization is an ongoing process where the secondary node is periodically syncs with the primary node. Synchronization is performed based on last synced timestamps of the primary and secondary nodes. Hence, there is no synchronization data queue in VMware Aria Operations.

**What is the actual witness node tolerance to missed polls?**

Witness node operations are not poll based. The witness node interacts only when one of the nodes is not able to communicate (after various checks) with nodes from the other fault domain.

**At what point in time will the primary node and primary replica node failover?**

The failover occurs only when the primary node is no longer accessible or not alive.

**When is the primary replica node promoted to the primary node?**

The primary replica node is promoted to the primary node in only two cases:

- When the existing primary node is down.
- The associated fault domain is down/offline.

**When the original primary node returns online, does it resume primary control? How does the data get synchronized?**

When operations return to normal, with both primary node and primary replica node online, the newly promoted primary node (formerly primary replica node) remains the new primary node and the new primary replica (formerly primary node) gets synced with the new primary node.

**What happens if connectivity between fault domains is completely interrupted, but then recovers?**

If communications between the fault domains is completely interrupted for several minutes, then one of the fault domains will automatically go offline. After the network interruption is restored, admin user needs to manually bring the fault domain online which will begin the data synchronization.

**What happens to the fault domains when the witness node becomes unavailable?**

While both fault domains are healthy and communicating with each other, the unavailability of the witness node will have no effect on the cluster; VMware Aria Operations will continue to function. If there is a communication issue between the fault domains, three situations could occur:

- Witness node is accessible from both fault domains – witness will bring one fault domain offline based on site health.
- Witness node is accessible from one fault domain only – The other fault domain will go offline automatically.
- Witness node is not accessible from both fault domains – Both fault domains will go offline.

**When the offline fault domain becomes available again, will the fault domains synchronize all data collected during the communication outage?**

The collected data is synchronized immediately once connectivity to the fault domain is restored and synchronized to capture all missed data.

**What happens when an analytics node is not able to communicate to analytics nodes in the other fault domain?**

If an analytics node is not able to communicate with all nodes from the other fault domain nor the witness node, it will go offline automatically. All nodes or entire fault domain that were taken offline automatically should be brought back online by the Admin user manually after ensuring that all communication issues have been resolved.

**If the maximum number of nodes in a standard cluster is 10 extra-large nodes, which supports 440,000 objects, why is the maximum number of nodes in continuous availability more with 12 extra-large nodes, which supports 264,000 objects?**

The 12 extra-large nodes are supported only in a continuous availability cluster and references a maximum of six extra large nodes across two separate fault domains. This permits an increase to the number of nodes over a standard cluster and allows for collection for a greater number of objects.

A possible design is six-large nodes in fault domain 1, and six extra-large nodes in fault domain 2, with a witness node in a third site. The latency requirements must be met such that latency between fault domain 1 and fault domain 2 is <10 ms. Details about latency, packet loss and bandwidth are listed in the Sizing Guidelines (KB 2093783).

**Is a load balancer supported with Continuous Availability?**

Yes, for more information about load balancer configuration, see VMware Aria Operations Load Balancing Configuration guide available under the **Related Resources** section of the VMware Aria Operations Documentation page.

**When the primary node is connected to the network again after a failover, what is the recommended procedure to return the original primary node to the primary role?**

It is not necessary to roll back the primary replica node to the primary node role or vice versa. If you still want to restore the old primary node to the primary role, then use “Take Node Offline/Online” on the new primary node or its fault domain (where the original primary node resides)

**Anytime a node goes offline or gets rebooted, is it necessary to bring the corresponding fault domain offline and then online to bring the node back online?**

All nodes, after reboot or bringing it offline/online, will automatically continue to work. No additional steps are necessary.

## Adapter and Management Packs Considerations

Adapters and management packs have specific configuration considerations.

### Normal Adapters

Normal adapters require a one-way communication to the monitored endpoint. Deploy normal adapters into collector groups, which are sized to handle a failover.

Following is a sample list of adapters provided by VMware for VMware Aria Operations. Additional adapters can be found on the VMware Solutions Exchange website.

- VMware vSphere
- Management Pack for NSX for vSphere
- Management Pack for VMware Integrated OpenStack
- Management Pack for Storage Devices
- Management Pack for Log Insight

## Hardware Requirements for Analytics Nodes, Witness Nodes, and Cloud Proxy

Analytics nodes, witness nodes, and cloud proxies have various hardware requirements for virtual machines and physical machines.

For information about the components to install on each server profile in your deployment, and the required hardware specifications, see the Sizing Guidelines (KB:2093783).

CPU requirements are 2.0 GHz minimum. 2.4 GHz is recommended. Storage requirements are based on the maximum supported resources for each node.

VMware Aria Operations has a high CPU requirement. In general, the more physical CPU that you assign to the analytics cluster, the better the performance. The cluster will perform better if the nodes stay within a single socket.

## Port Requirements for VMware Aria Operations

The most up-to-date technical information about ports can be found on Ports and Protocol.

## Small Deployment Profile for VMware Aria Operations

The small deployment profile contains a single large analytics node and is intended for systems that manage, for example, up to 20,000 resources. For accurate numbers, see the see the sizing guidelines.

### Virtual Appliance Name

The small deployment profile contains a single large analytics node, analytics-1.ra.local.

### Deployment Profile Support

For the configuration which the small deployment profile supports, see the sizing guidelines at KB2093783.

### Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytics-1.ra.local*

This is an example of a small deployment profile.

**Table 1: Adapter Properties**

| Collector Group | Collector | Adaptor | Resources |
|---|---|---|---|
| DEFAULT | analytics-1 | A | 2,000 |
| DEFAULT | analytics-1 | B | 4,000 |
| DEFAULT | analytics-1 | C | 2,000 |
| DEFAULT | analytics-1 | D | 3,000 |

## VMware Aria Operations Small Deployment Profile Architecture

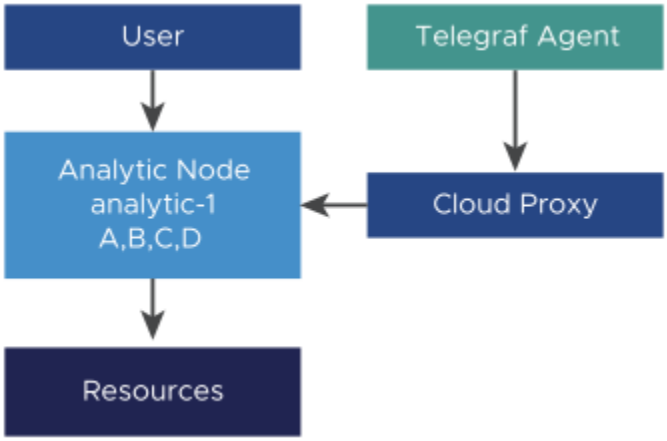


## Medium Deployment Profile for VMware Aria Operations

The medium deployment profile is intended for High Availability. For example, the medium deployment profile is intended for systems that manage 68,000 resources, 34,000 of which are activated for High Availability. For accurate numbers, see the sizing guidelines. In the medium deployment profile, adapters are deployed on the analytics nodes by default. If you experience problems with data ingestion, move these adapters to cloud proxies.

## Virtual Appliance Names

The medium deployment profile contains eight medium analytics nodes.

- analytics-1.ra.lcoal
- analytics-2.ra.lcoal
- analytics-3.ra.lcoal
- analytics-4.ra.lcoal
- analytics-5.ra.lcoal
- analytics-6.ra.lcoal
- analytics-7.ra.lcoal
- analytics-8.ra.lcoal

## Deployment Profile Support

For the configuration which the medium deployment profile supports, see the sizing guidelines at KB2093783.

## Load Balanced Addresses

- analytics.ra.local

## Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytics-1.ra.local*

This is an example of a medium deployment profile.

**Table 2: Adapter Properties**

| Collector Group | Collector | Adaptor | Resources |
|---|---|---|---|
| DEFAULT | analytics-1 | A | 2,000 |
| DEFAULT | analytics-2 | B | 4,000 |
| DEFAULT | analytics-3 | C | 2,000 |
| DEFAULT | analytics-4 | D | 3,000 |
| DEFAULT | analytics-5 | E | 1,000 |
| DEFAULT | analytics-6 | F | 2,000 |
| DEFAULT | analytics-7 | G | 1,500 |
| DEFAULT | analytics-8 | H | 4,500 |

## VMware Aria Operations Medium Deployment Profile Architecture

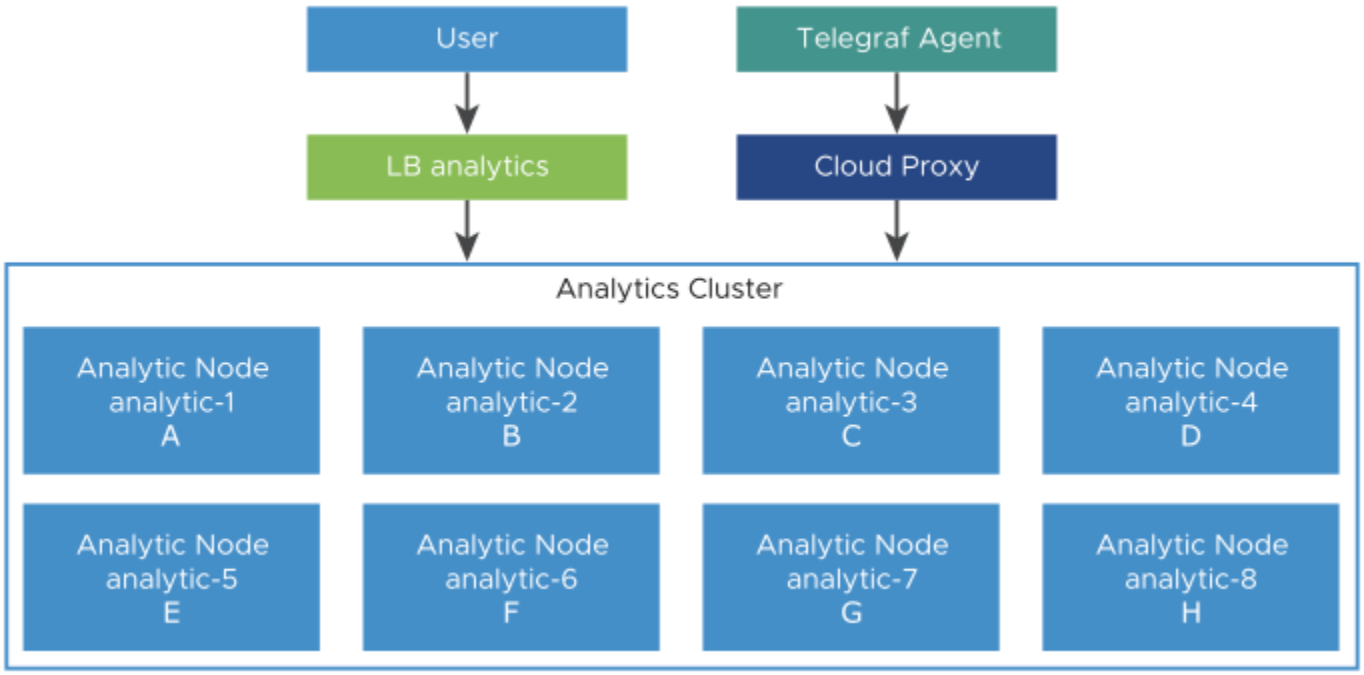


## Large Deployment Profile for VMware Aria Operations

The large deployment profile is intended for High Availability. For example, the large deployment profile is intended for systems that manage 128,000 resources, 64,000 of which are available with High Availability. For accurate numbers, see the sizing guidelines. All adapters are deployed to remote controllers in large deployment profiles to offload CPU usage from the analytics cluster.

## Virtual Appliance Names

The large deployment profile contains eight large analytics nodes, and large cloud proxies for adapters and Telegraf agents.

- analytics-1.ra.lcoal
- analytics-2.ra.lcoal
- analytics-3.ra.lcoal
- analytics-4.ra.lcoal
- analytics-5.ra.lcoal
- analytics-6.ra.lcoal
- analytics-7.ra.lcoal
- analytics-8.ra.lcoal

## Deployment Profile Support

For the configuration which the large deployment profile supports, see see the Sizing Guidelines (KB 2093783).

## Load Balanced Addresses

- `analytics.ra.local`

## Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *`analytics.refarch.local`*
- DNS Name = *`analytics-1.ra.local`* to DNS Name = *`analytics-8.ra.local`*
- DNS Name = *`remote-1.ra.local`* to DNS Name = *`remote-N.ra.local`*

This is an example of a large deployment profile.

**Table 3: Adapter Properties**

| Collector Group | Cloud Proxy | Adapter | Resources |
|---|---|---|---|
| 1 | CP-1 | A | 5,000 |
| 1 | CP-2 | B | 5,000 |
| | | Total | 10,000 |
| 2 | CP-3 | C | 10,000 |
| 2 | CP-4 | D | 5,000 |
| 2 | CP-5 | E | 5,000 |
| | | Total | 20,000 |

If a cloud proxy is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit for each cloud proxy.

## VMware Aria Operations Large Deployment Profile Architecture

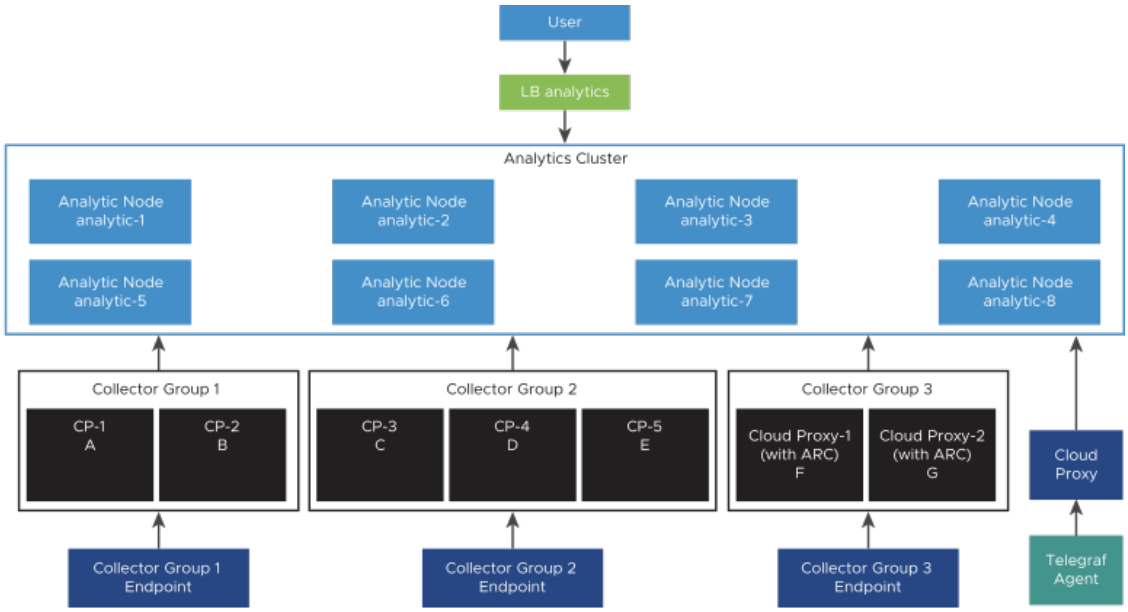


## Extra Large Deployment Profile for VMware Aria Operations

The extra-large deployment profile is intended Continuous Availability. For example, for systems that manage 240,000 resources, 120,000 of which are activated for Continuous Availability. For accurate numbers, see the sizing guidelines. This deployment is divided into two data centers and is the maximum supported analytics cluster deployment.

## Virtual Appliance Names

The extra-large deployment profile contains six extra-large analytics nodes. Large cloud proxies for adapters and witness node for continuous availability.

- `analytics-1.ra.local`
- `analytics-2.ra.local`
- `analytics-3.ra.local`
- `analytics-4.ra.local`
- `analytics-5.ra.local`
- `analytics-6.ra.local`

- witness-1.ra.local

## Deployment Profile Support

For the configuration which the extra large deployment profile supports, see the sizing guidelines at KB2093783.

## Load Balanced Addresses

- analytics.ra.local

## Certificate

The certificate must be signed by a Certificate Authority. The Subject Alternative Name contains the following information.

- DNS Name = *analytics.refarch.local*
- DNS Name = *analytics-1.ra.local* to *analytics-16.ra.local*
- DNS Name = *remote-1.ra.local* to *remote-N.ra.local*
- DNS Name = *witness-1.ra.local*

This is an example of an extra-large deployment profile. The adapter in the example provides N-1 redundancy, meaning, if two adapters support 20,000 resources, then a third adapter is added to attain a supported configuration that allows for a single failure.

**Table 4: Adapter Properties**

| Collector Group | Data Center | Cloud Proxy | Adapter | Resources |
|---|---|---|---|---|
| 1 | A | cp-1 | A | 5,000 |
| 1 | A | cp-2 | B | 5,000 |
| | | | Total | 10,000 |
| 2 | A | cp-3 | C | 2,000 |
| 2 | A | cp-3 | D | 2,000 |
| 2 | A | cp-3 | E | 1,000 |
| 2 | A | cp-4 | F | 7,000 |
| 2 | A | cp-5 | G | 8,000 |
| 2 | A | cp-6 | H | 5,000 |
| 2 | A | cp-7 | I | 6,000 |
| | | | Total | 31,000 |
| 3 | B | cp-8 | J | 10,000 |
| 3 | B | cp-9 | K | 5,000 |
| 3 | B | cp-10 | L | 5,000 |
| | | | Total | 20,000 |

If a cloud proxy is lost from these collector groups, you might have to manually rebalance the adapters to comply with the limit for each cloud proxy.

# VMware Aria Operations Extra Large Deployment Profile Architecture

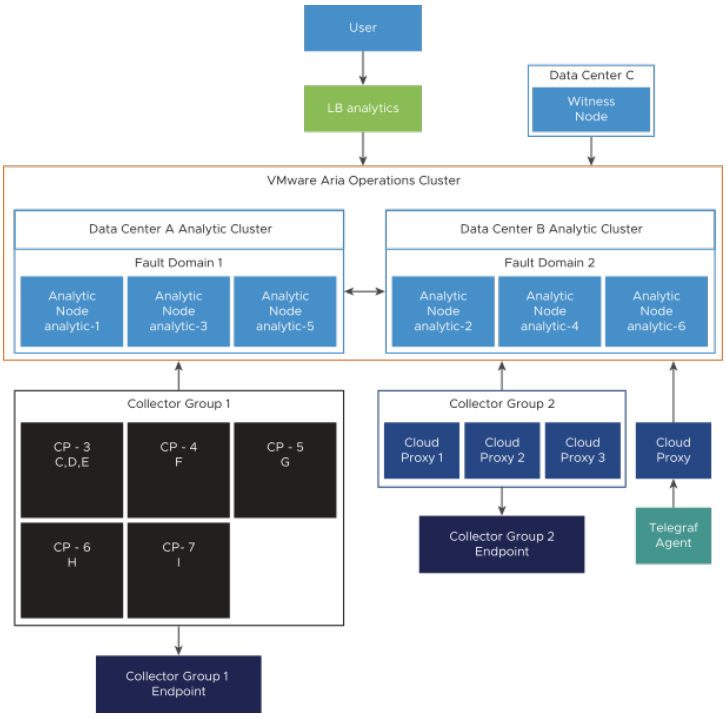

# VMware Aria Operations Secure Configuration Guide (8.18)

The documentation for *Secure Configuration* is intended to serve as a secure baseline for the deployment of VMware Aria Operations. Refer to this document when you are using system-monitoring tools to ensure that the secure baseline configuration is monitored and maintained for any unexpected changes on an ongoing basis.

Hardening activities that are not already set by default can be carried out manually.

## Intended Audience

This information is intended for administrators of VMware Aria Operations.

## VMware Aria Operations Security Posture

The security posture of VMware Aria Operations assumes a complete secure environment based on system and network configuration, organizational security policies, and best practices. It is important that you perform the hardening activities according to your organization's security policies and best practices.

The document is broken down into the following sections:

- Secure Deployment
- Secure Configuration
- Network Security
- Communication

The guide details the installation of the Virtual Application.

To ensure that your system is securely hardened, review the recommendations and assess them against your organization's security policies and risk exposure.

## Secure Deployment of VMware Aria Operations

You must verify the integrity of the installation media before you install the product to ensure authenticity of the downloaded files.

## Verify the Integrity of Installation Media

After you download the media, use the MD5/SHA1 sum value to verify the integrity of the download. Always verify the MD5/SHA1 hash after you download an OVA/OVF, offline bundle, or patch to ensure the integrity and authenticity of the downloaded files. If you obtain physical media from VMware and the security seal is broken, return the software to VMware for a replacement.

1. Compare the MD5/SHA1/SHA256 hash output with the value posted on the VMware website.

   SHA256, SHA1, or MD5 hash should match.

   **NOTE**
   The VMware Aria Operations`7.x-x.pak/8.x-x.pak` files are signed by the VMware software publishing certificate. VMware Aria Operations validates the signature of the PAK file before installation.

## How to Verify the Integrity of VMware Aria Operations Upgrade Pak Files Either from Trusted or Untrusted Sources

### Integrity Check from Trusted Sources

Every released or patch version of the downloadable product packages from VMware comes with its MD5 and SHA1 checksums in the VMware Customer Connect portal. The checksums can be used to verify if the downloaded file is intact, and in its original form. The above statement also applies to VMware Aria Operations installation PAK files (for cluster and cloud proxy upgrade, management packs, content packs, compliance packs, and so on).

**Integrity Check from Untrusted Sources**

If there is a lack of information about the download source, the digital signature of VMware Aria Operations installation PAK files can be verified manually before applying it. You can run the following steps to verify if the package contents have the correct signature by the trusted certificate.

1. Put the PAK file in `/storage/db/`. The directory of the primary node of the cluster.
   a) This can be achieved either by copying the PAK file to the target machine (using scp).
   b) Or by downloading the PAK file from the source (using wget).
2. Use the command below to verify the signature:

   ```
   python /usr/lib/vmwarevcopssuite/utilities/pakManager/bin/vcopsPakManager.py --action
   query_pak_signature -pak <path_to_pak_file>
   ```

   If the signature is valid, the following output will be printed:

   ```
   {
    "invalid_reason": null,
    "is_signature_valid": true,
    "is_signature_valid_certificate_untrusted": null,
    "is_signed": true,
    "pak_id": "PAK_NAME",
    "pak_version": "VERSION",
    "platform": [
    "Linux VA"
    ],
    "vcopssuiteinstall_build_number": null,
    "vcopssuiteinstall_build_type": null,
    "vcopssuitevm_build_number": "BUILD_NUMBER",
    "vcopssuitevm_build_type": "BUILD_TYPE"
   }
   ```

   Otherwise, in case of an invalid signature, the following output will be printed:

   ```
   {
    "invalid_reason": "MESSAGE",
    "is_signature_valid": false,
    "is_signature_valid_certificate_untrusted": null,
   ```

```
  "is_signed": false,
  "pak_id": "PAK_NAME",
  "pak_version": "VERSION",
  "platform": [
  "Windows",
  "Linux Non-VA",
  "Linux VA"
  ],
  "vcopssuiteinstall_build_number": null,
  "vcopssuiteinstall_build_type": null,
  "vcopssuitevm_build_number": null,
  "vcopssuitevm_build_type": null
}
```


## Hardening the Deployed Software Infrastructure

As part of your hardening process, you must harden the deployed software infrastructure that supports your VMware system.

Before you harden your VMware system, review and address security deficiencies in your supporting software infrastructure to create a completely hardened and secure environment. Software infrastructure elements to consider include operating system components, supporting software, and database software. Address security concerns in these and other components according to the manufacturer's recommendations and other relevant security protocols.

### Hardening the VMware vSphere Environment

VMware Aria Operations relies on a secure VMware vSphere environment to achieve the greatest benefits and a secured infrastructure.

Assess the VMware vSphere environment and verify that the appropriate level of vSphere hardening guidance is enforced and maintained.

For more guidance about hardening, see http://www.vmware.com/security/hardening-guides.html.

## Reviewing Installed and Unsupported Software

Vulnerabilities in unused software might increase the risk of unauthorized system access and disruption of availability. Review the software that is installed on VMware host machines and evaluate its use.

Do not install software that is not required for the secure operation of the system on any of the VMware Aria Operations node hosts. Uninstall unused or nonessential software.

Installing unsupported, untested, or unapproved software on infrastructure products such as VMware Aria Operations is a threat to the infrastructure.

To minimize the threat to the infrastructure, do not install or use any third-party software that is not supported by VMware on VMware supplied hosts.

Assess your VMware Aria Operations deployment and inventory of installed products to verify that no unsupported software is installed.

For more information about the support policies for third-party products, see the VMware support at http://www.vmware.com/security/hardening-guides.html.

#### Verify Third-Party Software

Do not use third-party software that VMware does not support. Verify that all third-party software is securely configured and patched in accordance with third-party vendor guidance.

Inauthentic, insecure, or unpatched vulnerabilities of third-party software installed on VMware host machines might put the system at risk of unauthorized access and disruption of availability. All software that VMware does not supply must be appropriately secured and patched.

If you must use third-party software that VMware does not support, consult the third-party vendor for secure configuration and patching requirements.

## VMware Security Advisories and Patches

VMware occasionally releases security advisories for products. Being aware of these advisories can ensure that you have the safest underlying product and that the product is not vulnerable to known threats. Assess the VMware Aria Operations installation, patching, and upgrade history and verify that the released VMware Security Advisories are followed and enforced.

It is recommended that you always remain on the most recent VMware Aria Operations release, as this will include the most recent security fixes also.

For more information about the current VMware security advisories, see http://www.vmware.com/security/advisories/.

## Secure Configuration of VMware Aria Operations

As a security best practice, you must secure the VMware Aria Operations console and manage Secure Shell (SSH), administrative accounts, and console access. Ensure that your system is deployed with secure transmission channels.

## Activating FIPS 140-2

FIPS 140-2 accreditation validates that an encryption solution meets a specific set of requirements designed to protect the cryptographic module from being cracked, altered, or otherwise tampered with. When FIPS 140-2 mode is activated, any secure communication to or from VMware Aria Operations 8.4 and above uses cryptographic algorithms or protocols that are allowed by the United States Federal Information Processing Standards (FIPS). FIPS mode turns on the cipher suites that comply with FIPS 140-2. Security related libraries that are shipped with VMware Aria Operations 8.4 and above are FIPS 140-2 certified. However, the FIPS 140-2 mode is not activated by default. FIPS 140-2 mode can be activated if there is a security compliance requirement to use FIPS certified cryptographic algorithms with the FIPS mode activated.

> **NOTE**
> Activating FIPS is a one-way action and cannot be deactivated after it is activated.

#### Activate FIPS during the initial cluster deployment

- Ensure a new deployment of a VMware Aria Operations cluster.
- Ensure that the Activate FIPS flag is appropriately used during the deployment of cluster nodes (OVF/OVA).

#### Activate FIPS on a working cluster

1. Navigate to `https://<VROPS IP>/admin/index.action`.
2. Login as an admin user.
3. Take the cluster offline to activate the Activate FIPS button in the **Administrator Settings** page.
4. Open the **Administrator Settings** tab in the left panel.
5. Click Activate FIPS under the **FIPS Setting** section.
6. Bring the cluster online.

**Verify that FIPS mode is Activated**
**From the Admin user interface:**

1. Navigate to `https://<VROPS IP>/admin/index.action`.
2. Login as the admin user.
3. Open the **Administrator Settings** tab from the left panel.
4. A **FIPS 140-2 Status** message appears.

## Activating Firewall Hardening

Activating firewall hardening restricts network access to internal services in VMware Aria Operations.

1. In a Web browser, navigate to the master node administration interface. `https://master-node-name-or-ip-address/admin`.
2. Enter the VMware Cloud Foundation Operations administrator username of admin.
3. Enter the VMware Cloud Foundation Operations administrator password and click **Log In**.
4. Click **Administrator Settings**, and then click **Security Settings** from the **Administrator Settings** page.
5. Click **Activate Firewall Hardening**.

## Secure the VMware Aria Operations Console

After you install VMware Aria Operations, you must log in for the first time and secure the console of each node in the cluster.

Install VMware Aria Operations.

1. Locate the node console in vCenter or by direct access.

   In vCenter, press Alt+F1 to access the login prompt. For security reasons, VMware Aria Operations remote terminal sessions are deactivated by default.

2. Log in as root.

   VMware Aria Operations does not allow you to access the command prompt until you create a root password.

3. At the prompt for a new password, enter the root password that you want and note it for future reference.
4. Reenter the root password.
5. Log out of the console.

## Change the Root Password

You can change the root password for any VMware Aria Operations primary or data node at any time by using the console.

Verify that the root password for the appliance meets your organization’s corporate password complexity requirements. If the account password starts with `$6$`, it uses a sha512 hash. This is the standard hash for all hardened appliances.

The root user passes the `pam_pwquality` module password complexity check, which is found in `/etc/pam.d/system-password`. All hardened appliances activate `enforce_for_root` for the `pam_pwquality` module, found in the `/etc/pam.d/system-password` file. The system remembers the last five passwords by default. Old passwords are stored for each user in the `/etc/security/opasswd` file.

1. Run the `# passwd` command at the root shell of the appliance.
2. To verify the hash of the root password, log in as root and run the `# more /etc/shadow` command.
   The hash information appears.
3. If the root password does not contain a sha512 hash, run the `passwd` command to change it.

**Manage Password Expiry**

Configure all account password expirations in accordance with your organization's security policies.

By default, the root account is set to a 365-day password expiry.
If the root password expires, you cannot reinstate it. You must implement site-specific policies to prevent administrative and root passwords from expiring.

1. Log in to your virtual appliance machines as root and run the `# more /etc/shadow` command to verify the password expiry on all accounts.
2. To modify the expiry of the root account, run the `# passwd -x 365 root` command.

   In this command, 365 specifies the number of days until password expiry. Use the same command to modify any user, substituting the specific account for **`root`** and replacing the number of days to meet the expiry standards of the organization.

   By default, the root password is set for 365 days.

## Managing Secure Shell, Administrative Accounts, and Console Access

For remote connections, all hardened appliances include the Secure Shell (SSH) protocol. SSH is deactivated by default on the hardened appliance.

SSH is an interactive command-line environment that supports remote connections to a VMware Aria Operations node. SSH requires high-privileged user account credentials. SSH activities generally bypass the role-based access control (RBAC) and audit controls of the VMware Aria Operations node.

As a best practice, deactivate SSH in a production environment and activate it only to diagnose or troubleshoot problems that you cannot resolve by other means. Leave it activated only while needed for a specific purpose and in accordance with your organization's security policies. If you activate SSH, ensure that it is protected against attack and that you activate it only for as long as required. Depending on your vSphere configuration, you can activate or deactivate SSH when you deploy your Open Virtualization Format (OVF) template.

As a simple test to determine whether SSH is activated on a machine, try to open a connection by using SSH. If the connection opens and requests credentials, then SSH is activated and is available for making connections.

**Secure Shell Root User**

Because VMware appliances do not include preconfigured default user accounts, the root account can use SSH to directly log in by default. Deactivate SSH as root as soon as possible.

To meet the compliance standards for nonrepudiation, the SSH server on all hardened appliances is preconfigured with the `AllowGroups wheel` entry to restrict SSH access to the secondary group wheel. For separation of duties, you can modify the `AllowGroups wheel` entry in the `/etc/ssh/sshd_config` file to use another group such as sshd.

The wheel group is activated with the `pam_wheel` module for superuser access, so members of the wheel group can use the su-root command, where the root password is required. Group separation activates users to use SSH to the appliance, but not to use the su command to log in as root. Do not remove or modify other entries in the AllowGroups field, which ensures proper appliance function. After making a change, restart the SSH daemon by running the `# service sshd restart` command.

### Activate or Deactivate Secure Shell on a VMware Aria Operations Node

You can activate Secure Shell (SSH) on a VMware Aria Operations node for troubleshooting. For example, to troubleshoot a server, you might require console access to the server through SSH. Deactivated SSH on a VMware Aria Operations node for normal operation.

1. Access the console of the VMware Aria Operations node from vCenter.
2. Select Login and press **Enter** to access the login prompt and then log in.
3. Run the `#systemctl is-enabled sshd` command.
4. If the sshd service is deactivated, run the `#systemctl enable sshd` command.
5. Run the `# systemctl start sshd` command to start the sshd service.
6. Run the `# systemctl stop sshd` command to stop the sshd service.

   You can also activate or deactivate Secure Shell from the **SSH Status** column of the VMware Aria Operations administration interface.

### Create a Local Administrative Account for Secure Shell

You must create local administrative accounts that can be used as Secure Shell (SSH) and that are members of the secondary `wheel` group, or both before you remove the root SSH access.

Before you deactivate direct root access, test that authorized administrators can access SSH by using **`AllowGroups`**, and that they can use the wheel group and the **`su`** command to log in as root.

1. Log in as root and run the following commands.

   ```
   # useradd username -d /home/vropsuser -g users -G wheel -m
   ```

   ```
   # passwd username
   ```

   Wheel is the group specified in **`AllowGroups`** for SSH access. To add multiple secondary groups, use `-G wheel,sshd`.

2. Switch to the user and provide a new password to ensure password complexity checking.

   ```
   # su – username
   ```

   ```
   username@hostname:~>passwd
   ```

   If the password complexity is met, the password updates. If the password complexity is not met, the password reverts to the original password, and you must rerun the password command.
   After you create the login accounts to allow SSH remote access and use the **`su`** command to log in as root using the wheel access, you can remove the root account from the SSH direct login.
3. To remove direct login to SSH, modify the `/etc/ssh/sshd_config` file by replacing `(#)PermitRootLogin yes` with `PermitRootLogin no`.

### Restrict Secure Shell Access

As part of your system hardening process, restrict Secure Shell (SSH) access by configuring the SSH package appropriately on all VMware virtual appliance host machines. Also maintain the required SSH key file permissions on these appliances.

1. Open the `/etc/ssh/sshd_config` file on your virtual appliance host machine in a text editor.
2. Change the generic entry for your production environment to include only the local host entries and the management network subnet for secure operations.

   Add the following line to the configuration file:
   ```
   AllowUsers root@127.0.0.1 root@::1 root@10.0.0.*
   ```

   In this example, all local host connections and connections that the clients make from the 10.0.0.0/24 subnet are allowed.
3. Save the file and close it.
4. Restart the SSH service by `systemctl restart sshd`.

### Maintain Secure Shell Key File Permissions

To maintain an appropriate level of security, configure Secure Shell (SSH) key file permissions.

1. View the public host key files, located in `/etc/ssh/*key.pub`.
2. Verify that these files are owned by root, that the group is owned by root, and that the files have permissions set to 0644.

   The permissions are (-rw-r--r--).
3. Close all files.
4. View the private host key files, located in `/etc/ssh/*key`.
5. Verify that root owns these files and the group, and that the files have permissions set to 0600.

   The permissions are (-rw-------).
6. Close all files.

### Harden the Secure Shell Server Configuration

Where possible, the Virtual Application Installation (OVF) has a default hardened configuration. Users can verify that their configuration is appropriately hardened by examining the server and client service in the global options section of the configuration file.

1. Open the `/etc/ssh/sshd_config` server configuration file and verify that the settings are correct.

| Setting | Status |
|---|---|
| Server Daemon Protocol | Protocol 2 |
| Ciphers | aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr |
| TCP Forwarding | AllowTCPForwarding no |

*Table continued on next page*

*Continued from previous page*

| Setting | Status |
|---|---|
| Server Gateway Ports | Gateway Ports no |
| X11 Forwarding | X11Forwarding no |
| SSH Service | Use the AllowGroups field and specify a group permitted to access and add members to the secondary group for users permitted to use the service. |
| GSSAPI Authentication | GSSAPIAuthentication no, if unused |
| Kerberos Authentication | KerberosAuthentication no, if unused |
| Local Variables (AcceptEnv global option) | Set to `disabled by commenting out` or `enabled for only LC_* or LANG variables` |
| Tunnel Configuration | PermitTunnel no |
| Network Sessions | MaxSessions 1 |
| Strict Mode Checking | Strict Modes yes |
| Privilege Separation | UsePrivilegeSeparation yes |
| rhosts RSA Authentication | RhostsRSAAuthentication no |
| Compression | Compression delayed or Compression no |
| Message Authentication code | hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-sha2-512,hmac-sha2-256 |
| User Access Restriction | PermitUserEnvironment no |
| KexAlgorithms | ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521 |

2. Ensure that the `ListenAddress` line is uncommented and set to a valid local IP.

   For example, `ListenAddress 0.0.0.0`
   Replace `0.0.0.0` with the IP address of the VMware Aria Operations node.

   For example, `ListenAddress 192.168.168.10`

3. Save your changes and close the file. At the command line, execute the following command to apply the changed settings: `# systemctl restart sshd.service`

### Harden the Secure Shell Client Configuration

As part of your system hardening monitoring process, verify hardening of the SSH client by examining the SSH client configuration file on virtual appliance host machines to ensure that it is configured according to VMware guidelines.

1. Open the SSH client configuration file, `/etc/ssh/ssh_config`, and verify that the settings in the global options section are correct.

| Setting | Status |
|---|---|
| Client Protocol | Protocol 2 |
| Client Gateway Ports | Gateway Ports no |
| GSSAPI Authentication | GSSAPIAuthentication no |
| Local Variables (SendEnv global option) | Provide only LC_* or LANG variables |
| Ciphers | aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr |

*Table continued on next page*

*Continued from previous page*

| Setting | Status |
|---|---|
| Message Authentication Codes | hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-sha2-512,hmac-sha2-256 |

2. Save your changes and close the file.

#### Deactivate SSH Access for the Admin User Account

As a security best practice, you can deactivate SSH access for the admin user account. The VMware Aria Operations admin account and the Linux admin account share the same password. Deactivating SSH access to the admin user enforces defense in depth by ensuring all users of SSH first login to a lesser privileged service account with a password that differs from the VMware Aria Operations admin account and then switch user to a higher privilege such as the admin or root.

1. Edit the `/etc/ssh/sshd_config` file.

   You can access this file from the command prompt.

2. Add the `DenyUsers admin` entry anywhere in the file and save the file.
3. To restart the sshd server, run the **`service sshd restart`** command.

## Set Boot Loader Authentication

To provide an appropriate level of security, configure boot loader authentication on your VMware virtual appliances. If the system boot loader requires no authentication, users with console access to the system might be able to alter the system boot configuration or boot the system to single user or maintenance mode, which can result in denial of service or unauthorized system access.

Because boot loader authentication is not set by default on the VMware virtual appliances, you must create a GRUB password to configure it.

1. Verify whether a boot password exists in the `/boot/grub/grub.cfg` file on your virtual appliances.
2. If no password exists, run the **`/usr/bin/grub2-mkpasswd-pbkdf2`** command on your virtual appliance.

   A password is generated, and the command supplies the hash output.

3. Add following lines at the end of `/etc/grub.d/40_custom`.

   ```
   set superusers="root"

   password_pbkdf2 root <hash of password>
   ```

4. Backup `/boot/grub/grub.cfg file` by using:

   ```
   cp /boot/grub/grub.cfg /boot/grub/grub.cfg.vropsbackup
   ```

5. Update the grub configuration by running the **`/usr/sbin/grub2-mkconfig -p /boot/grub/grub.cfg`** command.

> **NOTE**
> Important: Follow the upgrade procedure described below as otherwise, after upgrade, VMware Aria Operations will not start.

Upgrade procedure for VMware Aria Operations with a password protected boot loader.

1. Restore the old `grub.cfg` by running the following command:
   `cp /boot/grub/grub.cfg.vropsbackup /boot/grub/grub.cfg`
2. Upgrade VMware Aria Operations.
3. Follow all the steps described under **Set Boot Loader Authentication** after the upgrade of VMware Aria Operations.

## Monitor Minimal Necessary User Accounts

You must monitor existing user accounts and ensure that any unnecessary user accounts are removed.

1. Run the *host*`:~ # cat /etc/passwd` command and verify the minimal necessary user accounts:

```
root:x:0:0:root:/root:/bin/bash
bin:x:1:1:bin:/dev/null:/bin/false
daemon:x:6:6:Daemon User:/dev/null:/bin/false
messagebus:x:18:18:D-Bus Message Daemon User:/var/run/dbus:/bin/false
systemd-bus-proxy:x:72:72:systemd Bus Proxy:/:/bin/false
systemd-journal-gateway:x:73:73:systemd Journal Gateway:/:/bin/false
systemd-journal-remote:x:74:74:systemd Journal Remote:/:/bin/false
systemd-journal-upload:x:75:75:systemd Journal Upload:/:/bin/false
systemd-network:x:76:76:systemd Network Management:/:/bin/false
systemd-resolve:x:77:77:systemd Resolver:/:/bin/false
systemd-timesync:x:78:78:systemd Time Synchronization:/:/bin/false
nobody:x:65534:65533:Unprivileged User:/dev/null:/bin/false
apache:x:25:25:Apache Server:/srv/www:/bin/false
sshd:x:50:50:sshd PrivSep:/var/lib/sshd:/bin/false
ntp:x:87:87:Network Time Protocol:/var/lib/ntp:/bin/false
named:x:999:999::/var/lib/bind:/bin/false
admin:x:1000:1003::/home/admin:/bin/bash
postgres:x:1001:100::/var/vmware/vpostgres/14:/bin/bash
mpuser:x:1002:1003::/home/mpuser:/sbin/nologin
```

## Monitor Minimal Necessary Groups

You must monitor existing groups and members to ensure that any unnecessary groups or group access is removed.

1. Run the `<host>:~ # cat /etc/group` command to verify the minimum necessary groups and group membership.

```
root:x:0:admin
bin:x:1:daemon
sys:x:2:
kmem:x:3:
tape:x:4:
tty:x:5:
daemon:x:6:
floppy:x:7:
disk:x:8:
lp:x:9:
dialout:x:10:
audio:x:11:
video:x:12:
utmp:x:13:
usb:x:14:
cdrom:x:15:
adm:x:16:admin
messagebus:x:18:
systemd-journal:x:23:admin
input:x:24:
mail:x:34:
lock:x:54:
dip:x:30:
systemd-bus-proxy:x:72:
systemd-journal-gateway:x:73:
systemd-journal-remote:x:74:
systemd-journal-upload:x:75:
systemd-network:x:76:
systemd-resolve:x:77:
systemd-timesync:x:78:
nogroup:x:65533:
users:x:100:
```

```
sudo:x:27:

wheel:x:28:root,admin

apache:x:25:admin,apache

sshd:x:50:

ntp:x:87:

named:x:999:

admin:x:1003:

pivotal:x:1004:apache
```

## Resetting the VMware Aria Operations Administrator Password

As a security best practice, you can reset the VMware Aria Operations admin password for vApp installations.

1. Log in to the remote console of the primary node as root.
2. Enter the `$VMWARE_PYTHON_BIN $VCOPS_BASE/../vmware-vcopssuite/utilities/sliceConfiguration/bin/vcopsSetAdminPassword.py --reset` command and follow the prompts.

## Configure NTP on VMware Aria Operations

For critical time sourcing, deactivate host time synchronization and use the Network Time Protocol (NTP) on VMware appliances. You must configure a trusted remote NTP server for time synchronization. The NTP server must be an authoritative time server or at least synchronized with an authoritative time server.

The NTP daemon on VMware virtual appliances provides synchronized time services. NTP is deactivated by default, so you need to configure it manually. If possible, use NTP in production environments to track user actions and to detect potential malicious attacks and intrusions through accurate audit and log keeping. For information about NTP security notices, see the NTP Web site.
The NTP configuration file is located in the `/etc/ntp.conf` file on each appliance.

1. Navigate to the `/etc/ntp.conf` configuration file on your virtual appliance host machine.
2. Set the file ownership to `root:root`.
3. Set the permissions to `0640`.
4. To mitigate the risk of a denial-of-service amplification attack on the NTP service, open the `/etc/ntp.conf` file and ensure that the restrict lines appear in the file.

   ```
   restrict -4 default kod nomodify notrap nopeer noquery

   restrict -6 default kod nomodify notrap nopeer noquery

   restrict 127.0.0.1

   restrict -6 ::1
   ```

5. Save any changes and close the files.

   For information on NTP security notices, see http://support.ntp.org/bin/view/Main/SecurityNotice.

## Deactivate the TCP Timestamp Response on Linux

Use the TCP timestamp response to approximate the remote host's uptime and aid in further attacks. Additionally, some operating systems can be fingerprinted based on the behavior of their TCP time stamps.

1. Deactivate the TCP timestamp response on Linux.
   a) To set the value of `net.ipv4.tcp_timestamps` to `0`, run the `sysctl -w net.ipv4.tcp_timestamps=0` command.
   b) Add the `net.ipv4.tcp_timestamps=0` value in the default `sysctl.conf` file.

## TLS for Data in Transit

As a security best practice, ensure that the system is deployed with secure transmission channels.

### Configure Strong Protocols for VMware Aria Operations

Protocols such as SSLv2 and SSLv3 are no longer considered secure. In addition, TLS 1.0 and TLS 1.1 have also been deactivated and TLS 1.3 and TLS 1.2 are activated by default.

> **NOTE**
> When you upgrade from VMware Aria Operations 7.5 and above to 8.4 (and above), the user modifications to TLS settings are preserved. When you upgrade your VMware Aria Operations instance from 7.5 to 8.4 (and above), both TLS 1.0 and TLS 1.1 are deactivated on all the VMware Aria Operations nodes. TLS 1.3 and TLS 1.2 are the only protocols that are supported by default.

### Verify the Correct Use of Protocols in Apache HTTPD

VMware Aria Operations deactivates SSLv2, SSLv3, TLSv1, and TLSv1.1 by default. You must deactivate weak protocols on all load balancers before you put the system into production.

1. Run the **`grep SSLProtocol /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf | grep -v '#'`** command from the command prompt to verify that SSLv2, SSLv3, TLSv1, and TLSv1.1 are deactivated.

   If the protocols are deactivated, the command returns the following output: `SSLProtocol All -SSLv2 -SSLv3 -TLSv1 -TLSv1.1`.

2. To restart the Apache2 server, run the **`systemctl restart httpd`** command from the command prompt.

### Verify the Correct Use of Protocols in the GemFire TLS Handler

VMware Aria Operations deactivates SSLv3, TLS 1.0, and TLS 1.1 by default. You must deactivate weak protocols on all load balancers before you put the system into production.

1. Verify that the protocols are activated. To verify that the protocols are activated, run the following commands on each node:

   ```
   1. # grep inter_cluster.supported_protocols /storage/vcops/user/conf/ssl/secure-communications.properties

   or

   2. # grep default.supported_protocols /storage/vcops/user/conf/ssl/secure-communications.properties
   ```

If the result of command 1 is blank, that means that the inter_cluster properties are not specified directly and it uses default values which you can obtain by command 2.

2. Re-activate TLS 1.0 and TLS 1.1.
   a) Navigate to the administrator user interface to bring the cluster offline: `url/admin`.
   b) Click **Bring Offline**.
   c) To ensure that TLS 1.0 and TLS 1.1 are activated, run the following commands:

      If the result of command 1 is blank, use the following command:
      ```
      sed -i "/^[^#]*default.supported_protocols/ c\default.supported_protocols =
      TLSv1.2 TLSv1.1 TLSv1" /storage/vcops/user/conf/ssl/secure-
      communications.properties
      ```

      If the result of command 1 is not blank, use the following command:
      ```
      sed -i "/^[^#]*inter_cluster.supported_protocols/
      c\inter_cluster.supported_protocols = TLSv1.2 TLSv1.1 TLSv1" /storage/vcops/
      user/conf/ssl/secure-communications.properties
      ```

      Repeat this step for each node.

   d) Navigate to the administrator user interface to bring the cluster online.
   e) Click **Bring Online**.

### Configure VMware Aria Operations to Use Strong Ciphers

For maximum security, you must configure VMware Aria Operations components to use strong ciphers. To ensure that only strong ciphers are selected, deactivate the use of weak ciphers. Configure the server to support only strong ciphers and to use sufficiently large key sizes. Also, configure the ciphers in a suitable order.

VMware Aria Operations deactivates the use of cipher suites using the DHE key exchange by default. Ensure that you deactivate the same weak cipher suites on all load balancers before you put the system into production.

#### Using Strong Ciphers

The encryption cipher negotiated between the server and the browser determines the key exchange method and encryption strength that is used in a TLS session.

#### Verify the Correct Use of Cipher Suites in Apache HTTPD

For maximum security, verify the correct use of cipher suites in Apache httpd.

1. To verify the correct use of cipher suites in Apache httpd, run the **`grep SSLCipherSuite /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf | grep -v '#'`** command from the command prompt.

   If Apache httpd uses the correct cipher suites, the command returns the following output: `SSLCipherSuite HIGH:!aNULL!ADH:!EXP:!MD5:!3DES:!CAMELLIA:!PSK:!SRP:!DH:!AES256-GCM-SHA384:!AES256-SHA256:!AES256-SHA:!AES128-GCM-SHA256:!AES128-SHA256:!AES128-SHA:@STRENGTH`

2. To configure the correct use of cipher suites, run the **`sed -i "/^[^#]*SSLCipherSuite/ c\SSLCipherSuite HIGH:\!aNULL\!ADH:\!EXP:\!MD5:\!3DES:\!CAMELLIA:\!PSK:\!SRP:\!DH:\! AES256-GCM-SHA384:\!AES256-SHA256:\!AES256-SHA:\!AES128-GCM-SHA256:\!AES128-SHA256:\! AES128-SHA:@STRENGTH" /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf`** command from the command prompt.

   Run this command if the output in Step 1 is not as expected.

This command deactivates all cipher suites that use DH and DHE key exchange methods.

3. Run the **/etc/init.d/apache2 restart** command from the command prompt to restart the Apache2 server.
4. To reactivate DH, remove **!DH** from the cipher suites by running the **sed -i "/^[^#]*SSLCipherSuite/ c\SSLCipherSuite HIGH:\!aNULL\!ADH:\!EXP:\!MD5:\!3DES:\!CAMELLIA:\!PSK:\!SRP:\! AES256-GCM-SHA384:\!AES256-SHA256:\!AES256-SHA:\!AES128-GCM-SHA256:\!AES128-SHA256:\! AES128-SHA:@STRENGTH" /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf** command from the command prompt.
5. Run the **systemctl restart httpd** command from the command prompt to restart the Apache2 server.

### Verify the Correct Use of Cipher Suites in GemFire TLS Handler

For maximum security, verify the correct use of cipher suites in GemFire TLS Handler.

1. To verify that the cipher suites are activated, run the following commands on each node to verify that the protocols are activated:

   ```
   1. # grep inter_cluster.supported_cipher_suites /storage/vcops/user/conf/ssl/secure-communications.properties

   or

   2. # grep default.supported_cipher_suites /storage/vcops/user/conf/ssl/secure-communications.properties
   ```

   If the result of command 1 is blank, that means that the inter_cluster properties are not specified directly and it uses default values which you can obtain by command 2.

   The following result is expected:

   ```
   inter_cluster. supported_cipher_suites =

   TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,

   TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,

   TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,

   TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
   ```

   If the result of command 1 is blank, here is the expected result from command 2.

   ```
   default. supported_cipher_suites = TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,

   TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,

   TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,

   TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
   ```

2. Configure the correct cipher suites.
   a) Navigate to the administrator user interface at *URL*/admin.
   b) To bring the cluster offline, click **Bring Offline**.
   c) To configure the correct cipher suites, run the following commands:

      ```
      sed -i "/^[^#]*inter_cluster.supported_cipher_suites/
      c\inter_cluster.supported_cipher_suites =
      TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
      ```

```
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256" /storage/vcops/user/conf/ssl/secure-
communications.properties
```

If the result of command 1 is blank, use the following command to set cipher suites:

```
sed -i "/^[^#]*default.supported_cipher_suites/
c\default.supported_cipher_suites = TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256" /storage/vcops/user/conf/ssl/secure-
communications.properties
```

Repeat this step for each node.

d) Navigate to the administrator user interface at *URL*`/admin`.
e) Click **Bring Online**.

## Application Resources That Must be Protected

As a security best practice, ensure that the application resources are protected.

Follow the steps to ensure that the application resources are protected.

1. Run the `find / -path /proc -prune -o -type f -perm /6000 -ls` command to verify that the files have a well-defined SUID and GUID bits set.

   The following list appears:

```
141850     40 -rwsr-xr-x   1 root     root        40376 May 31 08:07 /usr/sbin/
unix_chkpwd

   143209     16 -rwsr-xr-x   1 root     root        15408 Feb 25  2021 /usr/sbin/
usernetctl

   142963     72 -rwsr-x---   1 root     root        66128 Oct 13  2022 /usr/libexec/
dbus-daemon-launch-helper

   141312    516 -rwsr-xr-x   1 root     root       524184 Aug  1 21:01 /usr/libexec/
ssh-keysign

   141930     60 -rwsr-xr-x   1 root     root        54464 Jun 21 16:27 /usr/bin/chsh

   141929     64 -rwsr-xr-x   1 root     root        60272 Jun 21 16:27 /usr/bin/chfn

   141927     56 -rwsr-xr-x   1 root     root        50384 Jun 21 16:27 /usr/bin/su

   140604     64 -rwsr-xr-x   1 root     root        61192 May 10  2022 /usr/bin/
mount

   142924     60 -rwsr-xr-x   1 root     root        53576 Feb 25  2021 /usr/bin/
crontab

   141938     60 -rwsr-xr-x   1 root     root        57000 Jun 21 16:27 /usr/bin/
newuidmap

   141926     76 -rwsr-xr-x   1 root     root        70088 Jun 21 16:27 /usr/bin/
passwd

   141928     80 -rwsr-xr-x   1 root     root        73984 Jun 21 16:27 /usr/bin/
chage

   141937     48 -rwsr-xr-x   1 root     root        46176 Jun 21 16:27 /usr/bin/
```

```
newgrp

   140621     36 -rwsr-xr-x   1 root     root        36224 May 10  2022 /usr/bin/
umount

   141458     36 -rwsr-xr-x   1 root     root        36248 Feb 24  2021 /usr/bin/
fusermount

   141936     60 -rwsr-xr-x   1 root     root        57008 Jun 21 16:27 /usr/bin/
newgidmap

   141934     92 -rwsr-xr-x   1 root     root        86720 Jun 21 16:27 /usr/bin/
gpasswd

   141931     32 -rwsr-xr-x   1 root     root        32376 Jun 21 16:27 /usr/bin/
expiry

   143459    272 -rwsr-xr-x   1 root     root       273600 Aug  4 03:02 /usr/bin/sudo
```

2. Run the `find / -path */proc -prune -o -nouser -print -o -nogroup -print` command to verify that all the files in the vApp have an owner.
   All the files have an owner if there are no results.
3. Run the `find / -name "*" -type f -not -path "*/sys*" -not -path "*/proc*" -not -path "*/dev*" -perm -o+w | xargs ls -lb` command to verify that none of the files are world writable files by reviewing permissions of all the files on the vApp.
   `Others` should not have write permission. The permissions on these files should be *`##4`* or *`##5`*, where *`#`* equals the default given set of permissions for the Owner and Group, such as *`6`* or *`7`*.
4. Run the `find / -path */proc -prune -o ! -user root -o -user admin -print` command to verify that the files are owned by the correct user.
   All the files belong to either `root` or `admin` if there are no results.
5. Run the `find /usr/lib/vmware-casa/ -type f -perm -o=w` command to ensure that files in the `/usr/lib/vmware-casa/` directory are not world writable.
   There must be no results.
6. Run the `find /usr/lib/vmware-vcops/ -type f -perm -o=w` command to ensure that files in the `/usr/lib/vmware-vcops/` directory are not world writable.
   There must be no results.
7. Run the `find /usr/lib/vmware-vcopssuite/ -type f -perm -o=w` command to ensure that files in the `/usr/lib/vmware-vcopssuite/` directory are not world writable.
   There must be no results.

## Apache Configuration

### Deactivate Web Directory Browsing

As a security best practice, ensure that a user cannot browse through a directory because it can increase the risk of exposure to directory traversal attacks.

1. Verify that web directory browsing is deactivated for all directories.
   a) Open the `/etc/httpd/httpd.conf` and `/usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf` files in a text editor.
   b) Verify that for each `<Directory>` listing, the option called **`Indexes`** for the relevant tag is omitted from the **`Options`** line.

### Verify Server Tokens for the Apache2 Server

As part of your system hardening process, verify server tokens for the Apache2 server. The Web server response header of an HTTP response can contain several fields of information. Information includes the requested HTML page, the Web server type and version, the operating system and version, and ports associated with the Web server. This information provides malicious users important information without the use of extensive tools.

The directive **`ServerTokens`** must be set to **`Prod`**. For example, **`ServerTokens Prod`**. This directive controls whether the response header field of the server that is sent back to clients includes a description of the operating system and information about compiled-in modules.

1. To verify server tokens, run the **`cat /etc/httpd/conf/extra/httpd-default.conf |grep ServerTokens`** command.
2. To modify **`ServerTokens Full`** to **`ServerTokens Prod`**, run the **`sed -i 's/\(ServerTokens\s\+\)Full/\1Prod/g' /etc/httpd/conf/extra/httpd-default.conf`** command.

### Deactivate the Trace Method for the Apache2 Server

In standard production operations, use of diagnostics can reveal undiscovered vulnerabilities that lead to compromised data. To prevent misuse of data, deactivate the HTTP `Trace` method.

1. To verify the `Trace` method for the Apache2 server, run the following command **`grep TraceEnable /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf`**.
2. To deactivate the `Trace` method for the Apache2 server, run the following command **`sed -i "/^[^#]*TraceEnable/ c\TraceEnable off" /usr/lib/vmware-vcopssuite/utilities/conf/vcops-apache.conf`**.

## Deactivate Configuration Modes

As a best practice, when you install, configure, or maintain VMware Aria Operations, you can modify the configuration or settings to activate troubleshooting and debugging of your installation.

Catalog and audit each of the changes you make to ensure that they are properly secured. Do not put the changes into production if you are not sure that your configuration changes are correctly secured.

## Managing Nonessential Software Components

To minimize security risks, remove or configure nonessential software from your VMware Aria Operations host machines.

Configure all software that you do not remove in accordance with manufacturer recommendations and security best practices to minimize the potential to create security breaches.

### Secure the USB Mass Storage Handler

Secure the USB mass storage handler to prevent it from loading by default on VMware Aria Operations appliances and to prevent its use as the USB device handler with the VMware Aria Operations appliances. Potential attackers can exploit this handler to install malicious software.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the `install usb-storage /bin/false` line appears in the file.
3. Save the file and close it.

### Secure the Bluetooth Protocol Handler

Secure the Bluetooth protocol handler on your VMware Aria Appliances to prevent potential attackers from exploiting it.

Binding the Bluetooth protocol to the network stack is unnecessary and can increase the attack surface of the host. Prevent the Bluetooth protocol handler module from loading by default on VMware Aria Appliances.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the line `install bluetooth /bin/false` appears in this file.
3. Save the file and close it.

### Secure the Stream Control Transmission Protocol

Prevent the Stream Control Transmission Protocol (SCTP) module from loading on VMware Aria appliances by default. Potential attackers can exploit this protocol to compromise your system.

Configure your system to prevent the SCTP module from loading unless it is absolutely necessary. SCTP is an unused IETF-standardized transport layer protocol. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the following line appears in this file.

   ```
   install sctp /bin/false
   ```

3. Save the file and close it.

### Secure the Datagram Congestion Control Protocol

As part of your system hardening activities, prevent the Datagram Congestion Control Protocol (DCCP) module from loading on VMware Aria appliances by default. Potential attackers can exploit this protocol to compromise your system.

Avoid loading the DCCP module, unless it is absolutely necessary. DCCP is a proposed transport layer protocol, which is not used. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes can cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the DCCP lines appear in the file.

   ```
   install dccp /bin/false
   ```

   ```
   install dccp_ipv4 /bin/false
   ```

   ```
   install dccp_ipv6 /bin/false
   ```

3. Save the file and close it.

### Secure Reliable Datagram Sockets Protocol

As part of your system hardening activities, prevent the Reliable Datagram Sockets (RDS) protocol from loading on your VMware Aria appliances by default. Potential attackers can exploit this protocol to compromise your system.

Binding the RDS protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the `install rds /bin/false` line appears in this file.
3. Save the file and close it.

### Secure the Transparent Inter-Process Communication Protocol

As part of your system hardening activities, prevent the Transparent Inter-Process Communication protocol (TIPC) from loading on your virtual appliance host machines by default. Potential attackers can exploit this protocol to compromise your system.

Binding the TIPC protocol to the network stack increases the attack surface of the host. Unprivileged local processes can cause the kernel to dynamically load a protocol handler by using the protocol to open a socket.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the `install tipc /bin/false` line appears in this file.
3. Save the file and close it.

### Secure Internet Packet Exchange Protocol

Prevent the Internetwork Packet Exchange (IPX) protocol from loading VMware Aria appliances by default. Potential attackers can exploit this protocol to compromise your system.

Avoid loading the IPX protocol module unless it is absolutely necessary. IPX protocol is an obsolete network-layer protocol. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the system to dynamically load a protocol handler by using the protocol to open a socket.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the line `install ipx /bin/false` appears in this file.
3. Save the file and close it.

### Secure AppleTalk Protocol

Prevent the AppleTalk protocol from loading on VMware Aria appliances by default. Potential attackers might exploit this protocol to compromise your system.

Avoid loading the AppleTalk Protocol module unless it is necessary. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes might cause the system to dynamically load a protocol handler by using the protocol to open a socket.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the line `install appletalk /bin/false` appears in this file.
3. Save the file and close it.

### Secure DECnet Protocol

Prevent the DECnet protocol from loading on your system by default. Potential attackers might exploit this protocol to compromise your system.

Avoid loading the DECnet Protocol module unless it is absolutely necessary. Binding this protocol to the network stack increases the attack surface of the host. Unprivileged local processes can cause the system to dynamically load a protocol handler by using the protocol to open a socket.

1. Open the DECnet Protocol `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the line `install decnet /bin/false` appears in this file.
3. Save the file and close it.

### Secure Firewire Module

Prevent the Firewire module from loading on VMware Aria appliances by default. Potential attackers might exploit this protocol to compromise your system.

Avoid loading the Firewire module unless it is necessary.

1. Open the `/etc/modprobe.d/modprobe.conf` file in a text editor.
2. Ensure that the line `install ieee1394 /bin/false` appears in this file.
3. Save the file and close it.

### Kernel Message Logging

The `kernel.printk` specification in the `/etc/sysctl.conf` file specifies the kernel print logging specifications.

There are 4 values specified:
- `console loglevel`. The lowest priority of messages printed to the console.
- `default loglevel`. The lowest level for messages without a specific log level.
- The lowest possible level for the console log level.
- The default value for console log level.

There are eight possible entries per value.
- `define KERN_EMERG "<0>" /* system is unusable */`
- `define KERN_ALERT "<1>" /* action must be taken immediately */`
- `define KERN_CRIT "<2>" /* critical conditions */`
- `define KERN_ERR "<3>" /* error conditions */`
- `define KERN_WARNING "<4>" /* warning conditions */`
- `define KERN_NOTICE "<5>" /* normal but significant condition */`
- `define KERN_INFO "<6>" /* informational */`
- `define KERN_DEBUG "<7>" /* debug-level messages */`

Set the `kernel.printk` values to `3 4 1 7` and ensure that the line `kernel.printk=3 4 1 7` exists in the `/etc/sysctl.conf` file.

## Additional Secure Configuration Activities

Block unnecessary ports on your host server that are not required.

### Deactivating Unnecessary Ports and Services

Verify the host server's firewall for the list of open ports that allow traffic.

Block all the ports that are not listed as a minimum requirement for VMware Aria Operations in the Configuring Ports and Protocols section of this document, or are not required. In addition, audit the services running on your host server and deactivate those that are not required.

## Network Security and Secure Communication

As a security best practice, review and edit the network communication settings of your VMware virtual appliances and host machines. You must also configure the minimum incoming and outgoing ports for VMware Aria Operations.

## Configuring Network Settings for Virtual Application Installation

To ensure that your VMware virtual appliance and host machines allow only safe and essential communication, review and edit their network communication settings.

### Set the Queue Size for TCP Backlog

As a security best practice, configure a default TCP backlog queue size on VMware appliance host machines. To mitigate TCP denial or service attacks, set an appropriate default size for the TCP backlog queue size. The recommended default setting is 1280.

1. Run the **`# cat /proc/sys/net/ipv4/tcp_max_syn_backlog`** command on each VMware appliance host machine.
2. Set the queue size for TCP backlog.
   a) Open the `/etc/sysctl.conf` file in a text editor.
   b) Set the default TCP backlog queue size by adding the following entry to the file.

      ```
      net.ipv4.tcp_max_syn_backlog=1280
      ```

   c) Save your changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Deny ICMPv4 Echoes to Broadcast Address

Responses to broadcast Internet Control Message Protocol (ICMP) echoes provide an attack vector for amplification attacks and can facilitate network mapping by malicious agents. Configuring your system to ignore ICMPv4 echoes provides protection against such attacks.

1. Run the `# cat /proc/sys/net/ipv4/icmp_echo_ignore_broadcasts` command to verify that the system is not sending responses to ICMP broadcast address echo requests.
2. Configure the host system to deny ICMPv4 broadcast address echo requests.
   a) Open the `/etc/sysctl.conf` file in a text editor.
   b) If the value for this entry is not set to `1`, add the `net.ipv4.icmp_echo_ignore_broadcasts=1` entry.
   c) Save the changes and close the file.
   d) Run **`# sysctl -p`** to apply the configuration.

### Configure the Host System to Deactivate IPv4 Proxy ARP

IPv4 Proxy ARP allows a system to send responses to ARP requests on one interface on behalf of hosts connected to another interface. You must deactivate IPv4 Proxy ARP to prevent unauthorized information sharing. Deactivate the setting to prevent leakage of addressing information between the attached network segments.

1. Run the **`# grep [01] /proc/sys/net/ipv4/conf/*/proxy_arp|egrep "default|all"`** command to verify whether the Proxy ARP is deactivated.
2. Configure the host system to deactivate IPv4 Proxy ARP.
   a) Open the `/etc/sysctl.conf` file in a text editor.
   b) If the values are not set to `0`, add the entries or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv4.conf.all.proxy_arp=0

      net.ipv4.conf.default.proxy_arp=0
      ```

   c) Save any changes you made and close the file.

d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Ignore IPv4 ICMP Redirect Messages

As a security best practice, verify that the host system ignores IPv4 Internet Control Message Protocol (ICMP) redirect messages. A malicious ICMP redirect message can allow a man-in-the-middle attack to occur. Routers use ICMP redirect messages to notify hosts that a more direct route exists for a destination. These messages modify the host's route table and are unauthenticated.

1. Run the **`# grep [01] /proc/sys/net/ipv4/conf/*/accept_redirects|egrep "default|all"`** command on the host system to check whether the host system ignores IPv4 redirect messages.
2. Configure the host system to ignore IPv4 ICMP redirect messages.
   a) Open the `/etc/sysctl.conf` file.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv4.conf.all.accept_redirects=0

      net.ipv4.conf.default.accept_redirects=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Ignore IPv6 ICMP Redirect Messages

As a security best practice, verify that the host system ignores IPv6 Internet Control Message Protocol (ICMP) redirect messages. A malicious ICMP redirect message might allow a man-in-the-middle attack to occur. Routers use ICMP redirect messages to tell hosts that a more direct route exists for a destination. These messages modify the host's route table and are unauthenticated.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/accept_redirects|egrep "default|all"`** command on the host system and check whether it ignores IPv6 redirect messages.
2. Configure the host system to ignore IPv6 ICMP redirect messages.
   a) Open the `/etc/sysctl.conf` to configure the host system to ignore the IPv6 redirect messages.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.accept_redirects=0

      net.ipv6.conf.default.accept_redirects=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv4 ICMP Redirects

As a security best practice, verify that the host system denies IPv4 Internet Control Message Protocol (ICMP) redirects. Routers use ICMP redirect messages to inform servers that a direct route exists for a particular destination. These messages contain information from the system's route table that might reveal portions of the network topology.

1. Run the **`# grep [01] /proc/sys/net/ipv4/conf/*/send_redirects|egrep "default|all"`** on the host system to verify whether it denies IPv4 ICMP redirects.
2. Configure the host system to deny IPv4 ICMP redirects.

a) Open the `/etc/sysctl.conf` file to configure the host system.
b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv4.conf.all.send_redirects=0

net.ipv4.conf.default.send_redirects=0
```

c) Save the changes and close the file.
d) Run `# sysctl -p` to apply the configuration.

**Configure the Host System to Log IPv4 Martian Packets**

As a security best practice, verify that the host system logs IPv4 Martian packets. Martian packets contain addresses that the system knows to be invalid. Configure the host system to log the messages so that you can identify misconfigurations or attacks in progress.

1. Run the **`# grep [01] /proc/sys/net/ipv4/conf/*/log_martians|egrep "default|all"`** command to check whether the host logs IPv4 Martian packets.
2. Configure the host system to log IPv4 Martian packets.
   a) Open the `/etc/sysctl.conf` file to configure the host system.
   b) If the values are not set to `1`, add the following entries to the file or update the existing entries accordingly. Set the value to `1`.

```
net.ipv4.conf.all.log_martians=1

net.ipv4.conf.default.log_martians=1
```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

**Configure the Host System to use IPv4 Reverse Path Filtering**

As a security best practice, configure your host machines to use IPv4 reverse path filtering. Reverse path filtering protects against spoofed source addresses by causing the system to discard packets with source addresses that have no route or if the route does not point towards the originating interface.

Configure your system to use reverse-path filtering whenever possible. Depending on the system role, reverse-path filtering might cause legitimate traffic to be discarded. In such cases, you might need to use a more permissive mode or deactivate reverse-path filtering altogether.

1. Run the **`# grep [01] /proc/sys/net/ipv4/conf/*/rp_filter|egrep "default|all"`** command on the host system to check whether the system uses IPv4 reverse path filtering.
2. Configure the host system to use IPv4 reverse path filtering.
   a) Open the `/etc/sysctl.conf` file to configure the host system.
   b) If the values are not set to `1`, add the following entries to the file or update the existing entries accordingly. Set the value to `1`.

```
net.ipv4.conf.all.rp_filter=1

net.ipv4.conf.default.rp_filter=1
```

   c) Save the changes and close the file.

d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv4 Forwarding

As a security best practice, verify that the host system denies IPv4 forwarding. If the system is configured for IP forwarding and is not a designated router, it can be used to bypass network security by providing a path for communication that is not filtered by network devices.

1. Run the **`# cat /proc/sys/net/ipv4/ip_forward`** command to verify whether the host denies IPv4 forwarding.
2. Configure the host system to deny IPv4 forwarding.
   a) Open the `/etc/sysctl.conf` to configure the host system.
   b) If the value is not set to `0`, add the following entry to the file or update the existing entry accordingly. Set the value to `0`.

      ```
      net.ipv4.ip_forward=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny Forwarding of IPv4 Source Routed Packets

Source-routed packets allow the source of the packet to suggest that routers forward the packet along a different path than what is configured on the router, which can be used to bypass network security measures.

This requirement applies only to the forwarding of source-routed traffic, such as when IPv4 forwarding is activated and the system is functioning as a router.

1. Run the **`# grep [01] /proc/sys/net/ipv4/conf/*/accept_source_route|egrep "default|all"`** command to verify whether the system does not use IPv4 source routed packets
2. Configure the host system to deny forwarding of IPv4 source routed packets.
   a) Open the `/etc/sysctl.conf` file with a text editor.
   b) If the values are not set to `0`, ensure that `net.ipv4.conf.all.accept_source_route=0` and the `net.ipv4.conf.default.accept_source_route=0` are set to `0`.
   c) Save and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Forwarding

As a security best practice, verify that the host system denies IPv6 forwarding. If the system is configured for IP forwarding and is not a designated router, it can be used to bypass network security by providing a path for communication that is not filtered by network devices.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/forwarding|egrep "default|all"`** command to verify whether the host denies IPv6 forwarding.
2. Configure the host system to deny IPv6 forwarding.
   a) Open the `/etc/sysctl.conf` to configure the host system.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.forwarding=0

      net.ipv6.conf.default.forwarding=0
      ```

c) Save the changes and close the file.
d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Use IPv4 TCP SYN Cookies

As a security best practice, verify that the host system uses IPv4 Transmission Control Protocol (TCP) SYN cookies. A TCP SYN flood attack might cause a denial of service by filling a system's TCP connection table with connections in the SYN_RCVD state. SYN cookies are used so as not to track a connection until a subsequent ACK is received, verifying that the initiator is attempting a valid connection and is not a flood source.

This technique does not operate in a fully standards-compliant manner, but is only activated when a flood condition is detected, and allows defense of the system while continuing to service valid requests.

1. Run the **`# cat /proc/sys/net/ipv4/tcp_syncookies`** command to verify whether the host system uses IPv4 TCP SYN cookies.
2. Configure the host system to use IPv4 TCP SYN cookies.
   a) Open the `/etc/sysctl.conf` to configure the host system.
   b) If the value is not set to `1`, add the following entry to the file or update the existing entry accordingly. Set the value to `1`.

      ```
      net.ipv4.tcp_syncookies=1
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Router Advertisements

As a security best practice, verify that the host system denies the acceptance of router advertisements and Internet Control Message Protocol (ICMP) redirects unless necessary. A feature of IPv6 is how systems can configure their networking devices by automatically using information from the network. From a security perspective, it is preferable to manually set important configuration information rather than accepting it from the network in an unauthenticated way.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra|egrep "default|all"`** command on the host system to verify whether the system denies the acceptance of router advertisements and ICMP redirects unless necessary.
2. Configure the host system to deny IPv6 router advertisements.
   a) Open the `/etc/sysctl.conf` file.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.accept_ra=0

      net.ipv6.conf.default.accept_ra=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Router Solicitations

As a security best practice, verify that host system denies IPv6 router solicitations unless necessary. The router solicitations setting determines how many router solicitations are sent when bringing up the interface. If addresses are assigned statically, there is no need to send any solicitations.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/router_solicitations|egrep "default|all"`** command to verify whether the host system denies IPv6 router solicitations unless necessary.
2. Configure the host system to deny IPv6 router solicitations.
   a) Open the `/etc/sysctl.conf`.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.router_solicitations=0

      net.ipv6.conf.default.router_solicitations=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Router Preference in Router Solicitations

As a security best practice, verify that your host system denies IPv6 router solicitations unless necessary. The router preference in the solicitations setting determines router preferences. If addresses are assigned statically, there is no need to receive any router preference for solicitations.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra_rtr_pref|egrep "default|all"`** on the host system to verify whether the host system denies IPv6 router solicitations.
2. Configure the host system to deny IPv6 router preference in router solicitations.
   a) Open the `/etc/sysctl.conf` file.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.accept_ra_rtr_pref=0

      net.ipv6.conf.default.accept_ra_rtr_pref=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Router Prefix

As a security best practice, verify that the host system denies IPv6 router prefix information unless necessary. The `accept ra pinfo` setting controls whether the system accepts prefix information from the router. If addresses are statically assigned, the system does not receive any router prefix information.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra_pinfo|egrep "default|all"`** to verify if that system denies IPv6 router prefix information.
2. Configure the host system to deny IPv6 router prefix.
   a) Open the `/etc/sysctl.conf` file.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.accept_ra_pinfo=0

      net.ipv6.conf.default.accept_ra_pinfo=0
      ```

   c) Save the changes and close the file.

d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Router Advertisement Hop Limit Settings

As a security best practice, verify that the host system denies IPv6 router advertisement Hop Limit settings from a router advertisement unless necessary. The **`accept_ra_defrtr`** setting controls whether the system accepts Hop Limit settings from a router advertisement. Setting it to `0` prevents a router from changing your default IPv6 Hop Limit for outgoing packets.

1. Run the `# grep [01] /proc/sys/net/ipv6/conf/*/accept_ra_defrtr|egrep "default|all"` command to verify that the host system denies IPv6 router Hop Limit settings.
2. If the values are not set to `0`, configure the host system to deny IPv6 router advertisement Hop Limit settings.
   a) Open the `/etc/sysctl.conf` file.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.accept_ra_defrtr=0

      net.ipv6.conf.default.accept_ra_defrtr=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Router Advertisement Autoconf Settings

As a security best practice, verify that the host system denies IPv6 router advertisement `autoconf` settings. The `autoconf` setting controls whether router advertisements can cause the system to assign a global unicast address to an interface.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/autoconf|egrep "default|all"`** command to verify whether the host system denies IPv6 router advertisement `autoconf` settings.
2. If the values are not set to `0`, configure the host system to deny IPv6 router advertisement autoconf settings.
   a) Open the `/etc/sysctl.conf` file.
   b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

      ```
      net.ipv6.conf.all.autoconf=0

      net.ipv6.conf.default.autoconf=0
      ```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

### Configure the Host System to Deny IPv6 Neighbor Solicitations

As a security best practice, verify that the host system denies IPv6 neighbor solicitations unless necessary. The `dad_transmits` setting determines how many neighbor solicitations are to be sent out per address including global and link-local, when you bring up an interface to ensure that the desired address is unique on the network.

1. Run the **`# grep [01] /proc/sys/net/ipv6/conf/*/dad_transmits|egrep "default|all"`** command to verify whether the host system denies IPv6 neighbor solicitations.
2. If the values are not set to `0`, configure the host system to deny IPv6 neighbor solicitations.

a) Open the `/etc/sysctl.conf` file.
b) If the values are not set to `0`, add the following entries to the file or update the existing entries accordingly. Set the value to `0`.

```
net.ipv6.conf.all.dad_transmits=0

net.ipv6.conf.default.dad_transmits=0
```

c) Save the changes and close the file.
d) Run `# sysctl -p` to apply the configuration.

#### Configure the Host System to Restrict IPv6 Maximum Addresses

As a security best practice, verify that the host restricts the maximum number of IPv6 addresses that can be assigned. The maximum addresses setting determines how many global unicast IPv6 addresses can be assigned to each interface. The default is 16 but you must set the number to the statically configured global addresses required.

1. Run the **`# grep [1] /proc/sys/net/ipv6/conf/*/max_addresses|egrep "default|all"`** command to verify whether the host system restricts the maximum number of IPv6 addresses that can be assigned.
2. If the values are not set to `1`, configure the host system to restrict the maximum number of IPv6 addresses that can be assigned.
   a) Open the `/etc/sysctl.conf` file.
   b) Add the following entries to the file or update the existing entries accordingly. Set the value to `1`.

```
net.ipv6.conf.all.max_addresses=1

net.ipv6.conf.default.max_addresses=1
```

   c) Save the changes and close the file.
   d) Run `# sysctl -p` to apply the configuration.

## Configuring Ports and Protocols

As a security best practice, deactivate all non-essential ports and protocols.

Configure the minimum incoming and outgoing ports for VMware Aria Operations components as required for important system components to operate in production.

#### Minimum Default Incoming Ports

As a security best practice, configure the incoming ports required for VMware Aria Operations to operate in production. The ports should be allowed/opened in local network for VMware Aria Operations inter-node communication and for customer to VMware Aria Operations communication.

The most up-to-date technical information for open ports can be found on Ports and Protocols.

## Cipher Suites and Protocols

The cipher suites and relevant protocols are listed when FIPS is in On and Off mode.

**NOTE**
It is strongly recommended that you do not use SSL, TLS 1.0, or TLS 1.1 protocols. TLS 1.2 and TLS 1.3 protocols should be considered as cornerstone configurations. Security of some of the cipher suites has degraded over time and as a result, some cipher suites are known to be insecure. Old or outdated cipher suites are often vulnerable to attacks. If they are used, the attacker may intercept or modify data in transit. It is recommended that you use only the following cipher suites:

| |
|---|
| TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 |
| TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 |
| TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 |
| TLS_AES_256_GCM_SHA384 |
| TLS_AES_128_GCM_SHA256 |

### Cipher Suites When FIPS is On

Here are the cipher suites lists when FIPS is On. The cipher suites are classified based on incoming, internode, and outbound connections. The cipher suite list is a comma-separated list.

#### Incoming Connections to VMware Aria Operations

**Table 5: Cipher Suites for Incoming Connections**

| Name | Cipher Suites |
|---|---|
| Configured Cipher Suites | |
| Apache Ciphers<br><br>Protocols - TLS 1.2, TLS 1.3 | ECDHE-RSA-AES256-GCM-SHA384<br>ECDHE-RSA-AES128-GCM-SHA256<br>ECDHE-RSA-AES256-SHA384<br>ECDHE-RSA-AES128-SHA256<br>ECDHE-RSA-AES256-SHA<br>ECDHE-RSA-AES128-SHA<br>AKE-AES256-GCM-SHA384<br>AKE-AES128-GCM-SHA256 |
| What you can configure: To find Apache relays to the OS cipher suite list, run the CLI command: `openssl ciphers -v 'HIGH:!aNULL!ADH:!EXP:!MD5:!3DES:!CAMELLIA:!PSK:!SRP:!DH:!AES256-GCM-SHA384:!AES256-SHA256:!AES256-SHA:!AES128-GCM-SHA256:!AES128-SHA256:!AES128-SHA:@STRENGTH'` | |

| Name | Cipher Suites |
|---|---|
| Configured Cipher Suites | |
| PostgreSQL Ciphers<br><br>Protocols - TLS 1.2, TLS 1.3 | TLS_AES_256_GCM_SHA384 |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_AES_128_GCM_SHA256<br>ECDHE-RSA-AES256-GCM-SHA384<br>ECDHE-RSA-AES128-GCM-SHA256<br>ECDHE-RSA-AES256-SHA384<br>ECDHE-RSA-AES128-SHA256 |
| What you can configure: To find PostgreSQL relays to the OS cipher suite list, run the CLI command: `openssl ciphers -v 'TLS_AES_256_GCM_SHA384:TLS_AES_128_GCM_SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES256-SHA384:ECDHE-ECDSA-AES256-SHA384:!aNULL:!eNULL:!EXPORT:!DES:!RC4:!3DES:!MD5:!PSK'` | |

**Internode Connections between VMware Aria Operations Nodes**

**Table 6: Cipher Suites for Internode Connections**

| Name | Cipher Suites |
|---|---|
| Configured Cipher Suites | |
| inter_cluster<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| What you can configure: | |
| All the possible cipher suites for internode connections. | TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_RSA_WITH_AES_256_CBC_SHA,<br>TLS_RSA_WITH_AES_128_CBC_SHA |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| **NOTE**<br>The PostgreSQL and Apache cipher suite lists must have an intersection with the inter_node cipher suite list. The inter_node proper cipher suite selection will avoid PostgreSQL and Apache from non-secure cipher suite usage. | |

**Outbound Connections from VMware Aria Operations**

Outbound cipher suites that are configured are classified into three types:

- Adapter to Source
- Authentication Sources
- Outbound Plugins

**Table 7: Adapter to Source**

| Name | Cipher Suites |
|---|---|
| All Adapters<br>Protocols - TLSv1.3, TLSv1.2, TLSv1.1, TLSv1 | TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

**Table 8: Authentication Sources**

| Name | Cipher Suites |
|---|---|
| vIDB<br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| vIDM<br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| sso_util<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| csp<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| LDAP<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

**Table 9: Outbound Plugins**

| Name | Cipher Suites |
|---|---|
| cprc_connection<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| marketplace_manager<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| email_sender<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| rest_sender<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| lint_rest_template<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

**Table 10: Outbound Cipher Suites that You Can Configure**

| Name | Cipher Suites |
|---|---|
| All the possible cipher suites you can configure for an outbound connection. | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

### Cipher Suites When FIPS is Off

Here are the lists of cipher suites when FIPS is Off. The cipher suites are classified based on incoming, internode, and outbound connections. The cipher suite list is a comma-separated list.

#### Incoming Connections to VMware Aria Operations

**Table 11: Cipher Suites for Incoming Connections**

| Name | Cipher Suites |
|---|---|
| Configured Cipher Suites | |
| Apache Ciphers<br><br>Protocols - TLS 1.2, TLS 1.3 | ECDHE-RSA-AES256-GCM-SHA384<br><br>ECDHE-RSA-CHACHA20-POLY1305<br><br>ECDHE-ARIA256-GCM-SHA384.<br><br>ECDHE-RSA-AES128-GCM-SHA256<br><br>ECDHE-ARIA128-GCM-SHA256<br><br>ECDHE-RSA-AES256-SHA384<br><br>ECDHE-RSA-AES128-SHA256<br><br>ECDHE-RSA-AES256-SHA<br><br>ECDHE-RSA-AES128-SHA<br><br>AES256-CCM8<br><br>AES256-CCM<br><br>ARIA256-GCM-SHA384<br><br>AES128-CCM8<br><br>AES128-CCM<br><br>ARIA128-GCM-SHA256 |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | AKE-AES256-GCM-SHA384<br>AKE-CHACHA20-POLY1305-SHA256<br>AKE-AES128-GCM-SHA256 |
| What you can configure: To find Apache relays to the OS cipher suite list, run the CLI command: `openssl ciphers -v 'HIGH:!aNULL!ADH:!EXP:!MD5:!3DES:!CAMELLIA:!PSK:!SRP:!DH:!AES256-GCM-SHA384:!AES256-SHA256:!AES256-SHA:!AES128-GCM-SHA256:!AES128-SHA256:!AES128-SHA:@STRENGTH'` | |

| Name | Cipher Suites |
|---|---|
| Configured Cipher Suites | |
| PostgreSQL Ciphers<br>Protocols - TLS 1.2 and TLS 1.3 | TLS_AES_256_GCM_SHA384<br>TLS_AES_128_GCM_SHA256<br>ECDHE-RSA-AES256-GCM-SHA384<br>ECDHE-RSA-AES128-GCM-SHA256<br>ECDHE-RSA-AES256-SHA384<br>ECDHE-RSA-AES128-SHA256 |
| What you can configure: To find PostgreSQL relays to the OS cipher suite list, run the CLI command: `openssl ciphers -v 'TLS_AES_256_GCM_SHA384:TLS_AES_128_GCM_SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES256-SHA384:ECDHE-ECDSA-AES256-SHA384:!aNULL:!eNULL:!EXPORT:!DES:!RC4:!3DES:!MD5:!PSK'` | |

## Internode Connections between VMware Aria Operations Nodes

**Table 12: Cipher Suites for Internode Connections**

| Name | Cipher Suites |
|---|---|
| Configured Cipher Suites | |
| inter_cluster<br>Protocols - TLSv1.3, TLSv1.2 | TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| What you can configure: | |
| All the possible cipher suites for internode connections. | TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDH_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDH_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA,<br>TLS_EMPTY_RENEGOTIATION_INFO_SCSV |
| **NOTE**<br>The PostgreSQL and Apache cipher suite lists must have an intersection with the inter_node cipher suite list. The inter_node proper cipher suite selection will avoid PostgreSQL and Apache from non-secure cipher suite usage. | |

**Outbound Connections from VMware Aria Operations**

Outbound cipher suites that are configured are classified into three types:

- Adapter to Source
- Authentication Sources
- Outbound Plugins

**Table 13: Adapter to Source**

| Name | Cipher Suites |
|---|---|
| All adapters<br>Protocols - TLSv1.3, TLSv1.2, TLSv1.1, TLSv1 | TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_AES_256_GCM_SHA384,<br>TLS_ECDH_RSA_WITH_AES_128_CBC_SHA,<br>TLS_CHACHA20_POLY1305_SHA256,<br>TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDH_RSA_WITH_AES_256_CBC_SHA,<br>TLS_RSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,<br>TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

**Table 14: Authentication Sources**

| Name | Cipher Suites |
|---|---|
| vIDB<br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,<br>TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_CBC_SHA256,<br>TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,<br>TLS_RSA_WITH_AES_256_CBC_SHA,<br>TLS_DHE_DSS_WITH_AES_128_GCM_SHA256,<br>TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,<br>TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,<br>TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA,<br>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,<br>TLS_RSA_WITH_AES_256_CBC_SHA256,<br>TLS_AES_128_GCM_SHA256, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| vIDM<br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| sso_util<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| csp<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| LDAP<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

**Table 15: Outbound Plugins**

| Name | Cipher Suites |
|---|---|
| cprc_connection<br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| marketplace_manager<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| email_sender<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| rest_sender<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |
| lint_rest_template<br><br>Protocols - TLSv1.3, TLSv1.2 | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

**Table 16: Outbound Cipher Suites that You Can Configure**

| Name | Cipher Suites |
|---|---|
| All the possible cipher suites you can configure for an outbound connection. | TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA256, TLS_DHE_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, |

*Table continued on next page*

*Continued from previous page*

| Name | Cipher Suites |
|---|---|
| | TLS_RSA_WITH_AES_256_CBC_SHA256, TLS_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA, TLS_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA, TLS_DHE_RSA_WITH_AES_128_CBC_SHA, TLS_CHACHA20_POLY1305_SHA256, TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256, TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384, TLS_DHE_RSA_WITH_AES_256_CBC_SHA256, TLS_EMPTY_RENEGOTIATION_INFO_SCSV, TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256, TLS_DHE_DSS_WITH_AES_256_CBC_SHA256, TLS_DHE_RSA_WITH_AES_256_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA, TLS_DHE_DSS_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 |

## Auditing and Logging on your VMware Aria Operations System

As a security best practice, set up auditing and logging on your VMware Aria Operations system.

The detailed implementation of auditing and logging is outside the scope of this document.

Remote logging to a central log host provides a secure store for logs. By collecting log files to a central host, you can easily monitor the environment with a single tool. You can also perform aggregate analysis and search for coordinated attacks on multiple entities within the infrastructure. Logging to a secure, centralized log server can help prevent log tampering and also provide a long-term audit record.

## Securing the Remote Logging Server

As a security best practice, ensure that the remote logging server can be configured only by an authorized user and is secure.

Attackers who breach the security of your host machine might search for and attempt to tamper with log files to cover their tracks and maintain control without being discovered.

## Use an Authorized NTP Server

Ensure that all the host systems use the same relative time source, including the relevant localization offset. You can correlate the relative time source to an agreed-upon time standard such as Coordinated Universal Time (UTC).

You can easily track and correlate an intruder's actions when you review the relevant log files. Incorrect time settings can make it difficult to inspect and correlate log files to detect attacks and can make auditing inaccurate. You can use at the least three NTP servers from outside time sources or configure a few local NTP servers on a trusted network that obtain their time from at least three outside time sources.

## Client Browser Considerations

As a security best practice, do not use VMware Aria Operations from untrusted or unpatched clients or from clients that use browser extensions.

# Getting Started with VMware Aria Operations (8.18)

The *Getting Started with VMware Aria Operations Guide* provides information about deploying the VMware®VMware Aria Operations virtual appliance, including how to create and configure the VMware Aria Operations cluster. The VMware Aria Operations installation process consists of deploying the VMware Aria Operations virtual appliance once for each cluster node, and accessing the product to finish setting up the application.

## Intended Audience

This information is intended for anyone who wants to install and configure VMware Aria Operations by using a virtual appliance deployment. The information is written for experienced virtual machine administrators who are familiar with enterprise management applications and data center operations

For administrators who want to deploy the VMware Aria Operations virtual appliance programmatically, the VMware Aria Operations CaSA API documentation is available in HTML format and is installed with your VMware Aria Operations instance. For example, if the URL of your instance is `https://operations.example.com`, the API reference is available from `https://operations.example.com/casa/api-guide.html`.

## About Installing

You prepare for VMware Aria Operations installation by evaluating your environment and deploying enough VMware Aria Operations cluster nodes to support how you want to use the product.

## Workflow of VMware Aria Operations Installation

The VMware Aria Operations virtual appliance installation process consists of deploying the VMware Aria Operations OVA, once for each cluster node, accessing the product to set up cluster nodes according to their role, and logging in to configure the installation.

**Figure 1: VMware Aria Operations Installation Architecture**

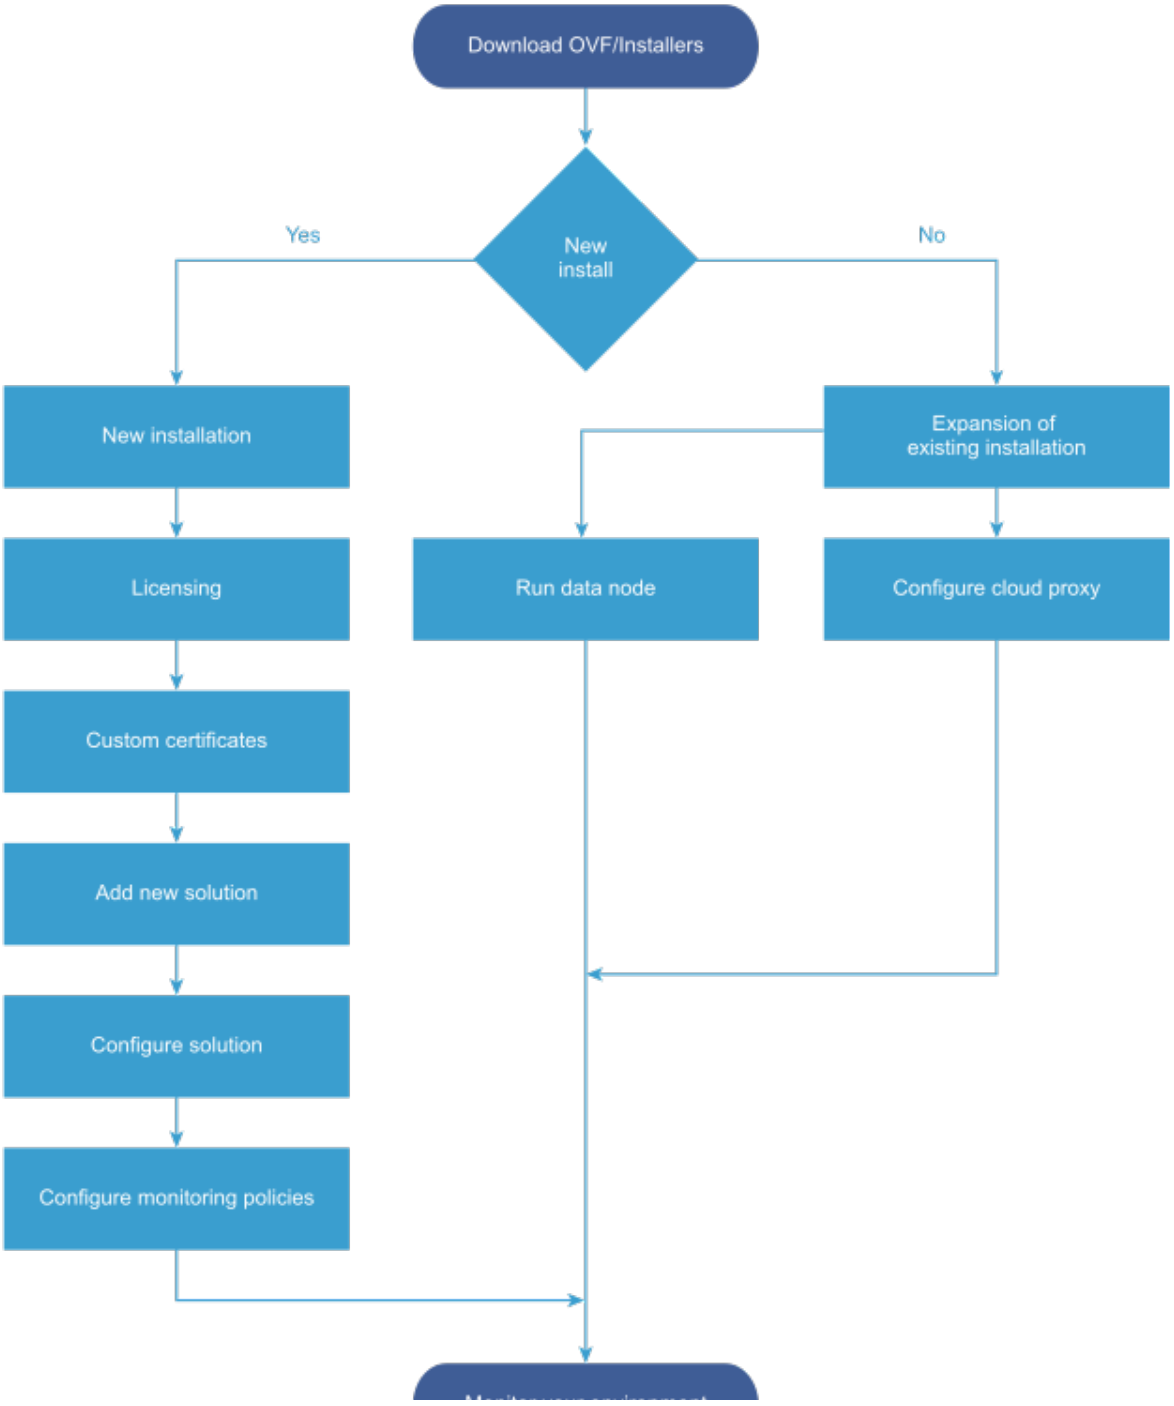


To automate installation, configuration, upgrade, patch, configuration management, drift remediation and health from within a single pane of glass, you can use Fleet Management. If you are a new user, click here to install Fleet Management. This provides the IT Managers of Cloud admin resources to focus on business-critical initiatives, while improving time to value (TTV),reliability, and and consistency.

You can also install upgrade VMware Aria Operations by using Fleet Management

# Sizing the VMware Aria Operations Cluster

Sizing the Cluster
The resources needed for VMware Aria Operations depend on how large of an environment you expect to monitor and analyze, how many metrics you plan to collect, and how long you need to store the data.

It is difficult to broadly predict the CPU, memory, and disk requirements that will meet the needs of a particular environment. There are many variables, such as the number and type of objects collected, which includes the number and type of adapters installed, the presence of HA, the duration of data retention, and the quantity of specific data points of interest, such as symptoms, changes, and so on.

VMware expects VMware Aria Operations sizing information to evolve, and maintains Knowledge Base articles so that sizing calculations can be adjusted to adapt to usage data and changes in versions of VMware Aria Operations.

Knowledge Base article 2093783

The Knowledge Base articles include overall maximums, plus spreadsheet calculators in which you enter the number of objects and metrics that you expect to monitor. To obtain the numbers, some users take the following high-level approach, which uses VMware Aria Operations itself.

1. Review this guide to understand how to deploy and configure a VMware Aria Operations node.
2. Deploy a temporary VMware Aria Operations node.
3. Configure one or more adapters, and allow the temporary node to collect overnight.
4. Access the Cluster Management page on the temporary node.
5. Using the Adapter Instances list in the lower portion of the display as a reference, enter object and metric totals of the different adapter types into the appropriate sizing spreadsheet from Knowledge Base article 2093783.
6. Deploy the VMware Aria Operations cluster based on the spreadsheet sizing recommendation. You can build the cluster by adding resources and data nodes to the temporary node or by starting over.

If you have a large number of adapters, you might need to reset and repeat the process on the temporary node until you have all the totals you need. The temporary node will not have enough capacity to simultaneously run every connection from a large enterprise.

Another approach to sizing is through self monitoring. Deploy the cluster based on your best estimate, but create an alert for when capacity falls below a threshold, one that allows enough time to add nodes or disk to the cluster. You also have the option to create an email notification when thresholds are passed.

During internal testing, a single-node vApp deployment of VMware Aria Operations that monitored 8,000 virtual machines ran out of disk storage within one week.

### Add Data Disk Space to a VMware Aria Operations vApp Node

Add Data Disk Space to a vApp Node
You add to the data disk of VMware Aria Operations vApp nodes when space for storing the collected data runs low.

- Note the disk size of the analytics cluster nodes. When adding disk, you must maintain uniform size across analytics cluster nodes. For more information, see KB article 318768.
- Use the VMware Aria Operations administration interface to take the node offline.
- Verify that you are connected to a vCenter system with a vSphere Client, and log in to the vSphere Client.

1. Shut down the virtual machine for the node.
2. Edit the hardware settings of the virtual machine, and add another disk.

   **NOTE**
   Do not expand disks. VMware Aria Operations does not support expanding disks.

3. Power on the virtual machine for the node.

During the power-on process, the virtual machine expands the VMware Aria Operations data partition.

**Add Data Disk Space to a VMware Aria Operations Linux Node**

Add Data Disk Space to a Linux Node
You add to the data disk of VMware Aria Operations Linux nodes when space for storing the collected data runs low.

Note the disk size of the analytics cluster nodes. When adding disk, you must maintain uniform size across analytics cluster nodes.

The following example is for a Linux system.

1. Add a new disk to the system, and partition and format the disk as needed.
2. Use the VMware Aria Operations administration interface to take the cluster offline.
3. Stop the `vmware-casa` service.
4. Move the contents of `/storage/db` into a directory on the new disk.
5. Create a symbolic link from the new directory back to `/storage/db`, so that `/storage/db` now references the new disk.
6. Start the `vmware-casa` service.
7. Bring the cluster online.

## Complexity of Your Environment

When you deploy VMware Aria Operations, the number and nature of the objects that you want to monitor might be complex enough to recommend a Professional Services engagement.

**Complexity Levels**

Every enterprise is different in terms of the systems that are present and the level of experience of deployment personnel. The following table presents a color-coded guide to help you determine where you are on the complexity scale.

- Green
  Your installation only includes conditions that most users can understand and work with, without assistance. Continue your deployment.
- Yellow
  Your installation includes conditions that might justify help with your deployment, depending on your level of experience. Consult your account representative before proceeding, and discuss using Professional Services.
- Red
  Your installation includes conditions that strongly recommend a Professional Services engagement. Consult your account representative before proceeding, and discuss using Professional Services.

Note that these color-coded levels are not firm rules. Your product experience, which increases as you work with VMware Aria Operations and in partnership with Professional Services, must be taken into account when deploying VMware Aria Operations.

**Table 17: Effect of Deployment Conditions on Complexity**

| Complexity Level | Current or New Deployment Condition | Additional Notes |
|---|---|---|
| Green | You run only one VMware Aria Operations deployment. | Lone instances are usually easy to create in VMware Aria Operations. |

*Table continued on next page*

*Continued from previous page*

| Complexity Level | Current or New Deployment Condition | Additional Notes |
|---|---|---|
| Green | Your deployment includes a management pack that is listed as Green according to the compatibility guide on the VMware Solutions Exchange Web site. | The compatibility guide indicates whether the supported management pack for VMware Aria Operations is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected.<br><br>Note that the terms *solution*, *management pack*, *adapter*, and *plug-in* are used somewhat interchangeably. |
| Yellow | You run multiple instances of VMware Aria Operations. | Multiple instances are typically used to address scaling or operator use patterns. |
| Yellow | Your deployment includes a management pack that is listed as Yellow according to the compatibility guide on the VMware Solutions Exchange Web site. | The compatibility guide indicates whether the supported management pack for VMware Aria Operations is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected. |
| Yellow | You are deploying a multiple-node VMware Aria Operations cluster. | Multiple nodes are typically used for scaling out the monitoring capability of VMware Aria Operations. |
| Yellow | Your new VMware Aria Operations instance will include a Linux based deployment. | Linux deployments are not as common as vApp deployments and often need special consideration. |
| Yellow | Your VMware Aria Operations instance will use high availability (HA). | High availability and its node failover capability is a unique multiple-node feature that you might want additional help in understanding. |
| Yellow | You want help in understanding the new or changed features in VMware Aria Operations and how to use them in your environment. | VMware Aria Operations is different than vCenter Operations Manager in areas such as policies, alerts, compliance, custom reporting, or badges. In addition, VMware Aria Operations uses one consolidated interface. |
| Red | You run multiple instances of VMware Aria Operations, where at least one | Multiple instances are typically used to address scaling, operator use patterns, |

*Table continued on next page*

*Continued from previous page*

| Complexity Level | Current or New Deployment Condition | Additional Notes |
|---|---|---|
| | includes virtual desktop infrastructure (VDI). | or because separate VDI (V4V monitoring) and non-VDI instances are needed. |
| Red | Your deployment includes a management pack that is listed as Red according to the compatibility guide on the VMware Solutions Exchange Web site. | The compatibility guide indicates whether the supported management pack for VMware Aria Operations is a compatible 5.x one or a new one designed for this release. In some cases, both might work but produce different results. Regardless, users might need help in adjusting their configuration so that associated data, dashboards, alerts, and so on appear as expected. |
| Red | You are deploying multiple VMware Aria Operations clusters. | Multiple clusters are typically used to isolate business operations or functions. |
| Red | Your current VMware Aria Operations deployment required a Professional Services engagement to install it. | If your environment was complex enough to justify a Professional Services engagement in the previous version, it is possible that the same conditions still apply and might warrant a similar engagement for this version. |
| Red | Professional Services customized your VMware Aria Operations deployment. Examples of customization include special integrations, scripting, nonstandard configurations, multiple level alerting, or custom reporting. | If your environment was complex enough to justify a Professional Services engagement in the previous version, it is possible that the same conditions still apply and might warrant a similar engagement for this version. |

## About VMware Aria Operations Cluster Nodes

Cluster Nodes
All VMware Aria Operations clusters consist of a primary node, an optional replica node for high availability or continuously availability, and optional data nodes.

When you install VMware Aria Operations, you use a VMware Aria Operations vApp deployment to create role-less nodes. After the nodes are created and have their names and IP addresses, you use an administration interface to configure them according to their role.

You can create role-less nodes all at once or as needed. A common as-needed practice might be to add nodes to scale out VMware Aria Operations to monitor an environment as the environment extends larger.

The following node types make up the VMware Aria Operations analytics cluster:

**Primary Node**

The primary node is the initial, required node in VMware Aria Operations. All other nodes are managed by the primary node.

In a single-node installation, the primary node manages itself, has adapters installed on it, and performs all data collection and analysis.

**Data Node**

In larger deployments, additional data nodes have adapters installed and perform collection and analysis.

Larger deployments usually include adapters only on the data nodes so that primary and replica node resources can be dedicated to cluster management.

**Replica Node**

To use VMware Aria Operations high availability (HA) and continuous availability (CA) the cluster requires that you convert a data node into a replica of the primary node.

The following node types are a member of the VMware Aria Operations cluster but not part of the analytics cluster:

**Witness Node**

To use VMware Aria Operations continuous availability (CA), the cluster requires that you have a witness node. Each VMware Aria Operations cluster can have only one witness node. If the network connection between the two fault domains is lost, the witness node acts as a decision maker regarding the availability of VMware Aria Operations.

**NOTE**
If you require an agent to collect data, you must deploy a cloud proxy. For more information on how to deploy a cloud proxy, see Installing .

## About VMware Aria Operations High Availability

About High Availability
VMware Aria Operations supports high availability (HA). HA creates a replica for the VMware Aria Operations primary node and protects the analytics cluster against the loss of a node.

With HA, data stored in the primary node is always 100% backed up on the replica node. To activate HA, you must have at least one data node deployed, in addition to the primary node. If you have more than one data node, the data stored in the primary node can be stored and replicated in any of the other nodes. But in case the primary node fails, only the replica node can function as the replacement of the primary node.

- HA is not a disaster recovery mechanism. HA protects the analytics cluster against the loss of only one node, and because only one loss is supported, you cannot stretch nodes across vSphere clusters in an attempt to isolate nodes or build failure zones.
- When HA is activated, the replica can take over all functions that the primary provides, were the primary to fail for any reason. If the primary fails, failover to the replica is automatic and requires only two to three minutes of VMware Aria Operations downtime to resume operations and restart data collection.
  When a primary node problem causes failover, the replica node becomes the primary node, and the cluster runs in degraded mode. To get out of degraded mode, take one of the following steps.
  - Return to HA mode by correcting the problem with the primary node. When a primary node exits an HA-activated cluster, primary node does not rejoin with the cluster without manual intervention. Therefore, restart the VMware Aria Operations Analytics process on the downed node to change its role to replica and rejoin the cluster.
  - Remove the failed primary node then re-activate HA by converting a data node into replica. Removed primary nodes cannot be repaired and readded to VMware Aria Operations.
  - Remove the old, failed primary node and then change to non-HA operation by deactivating HA. Removed primary nodes cannot be repaired and readded to VMware Aria Operations.
- In the administration interface, after an HA replica node takes over and becomes the new primary node, you cannot remove the previous, offline primary node from the cluster. In addition, the previous node remains listed as a primary node. To refresh the display and activate removal of the node, refresh the browser.
- When HA is activated, the cluster can survive the loss of one data node without losing any data. However, HA protects against the loss of only one node at a time, of any kind, so simultaneously losing data and primary/replica nodes, or two or more data nodes, is not supported. Instead, VMware Aria Operations HA provides additional application level data protection to ensure application level availability.

- When HA is activated, it lowers VMware Aria Operations capacity and processing by half, because HA creates a redundant copy of data throughout the cluster, and the replica backup of the primary node. Consider your potential use of HA when planning the number and size of your VMware Aria Operations cluster nodes. See Sizing the Cluster.
- When HA is activated, deploy analytics cluster nodes on separate hosts for redundancy and isolation. One option is to use anti-affinity rules that keep nodes on specific hosts in the vSphere cluster.
  If you cannot keep the nodes separate, you should not activate HA. A host fault might cause the loss of more than one node, which is not supported, and all of VMware Aria Operations can become unavailable.

  The opposite is also true. Without HA, you can keep nodes on the same host, and it will not make a difference. Without HA, the loss of even one node can make all of VMware Aria Operations unavailable.
- When you power off the data node and change the network settings of the VM, this affects the IP address of the data node. After this point, the HA cluster is no longer accessible and all the nodes have a status of "Waiting for analytics". Verify that you have used a static IP address.
- When you remove a node that has one or more vCenter adapters configured to collect data from a HA-activated cluster, one or more vCenter adapters associated with that node stops collecting. You change the adapter configuration to pin them to another node before removing the node.
- Administration UI shows the resource cache count, which is created for active objects only, but the Inventory displays all objects. Therefore, when you remove a node from a HA-activated cluster allowing the vCenter adapters collect data and rebalance each node, the Inventory displays a different quantity of objects from that shown in the Administration UI.

## About VMware Aria Operations Continuous Availability

VMware Aria Operations supports continuous availability (CA). CA separates the VMware Aria Operations cluster into two fault domains, stretching across vSphere clusters, and protects the analytics cluster against the loss of an entire fault domain.

You can configure the analytics cluster with Continuous Availability. This allows the cluster nodes to be stretch across two fault-domains. A fault domain consists of one or more analytics nodes grouped according to their physical location in the data center. With CA, the two fault domains permit VMware Aria Operations to tolerate failures of an entire physical location and failures from resources dedicated to a single fault domain.

To activate continuous availability within VMware Aria Operations, the witness node must be deployed in the cluster. The VMware Aria Operations cluster can have only one witness node. The witness node does not collect nor store data. In a situation where network connectivity the two fault-domains is lost, the cluster would go into a split-brain situation. This situation is detected by the Witness Node and one of the fault domains will go offline to avoid data inconsistency issues. You will see a **Bring Online** button on the admin UI of the nodes which are made offline by the witness node. Before using this option to bring the fault domain online, ensure that the network connectivity between the nodes across the two fault domains is restored and stable. Once confirmed you can bring the fault domain online.

With CA, the data stored in the primary node and data nodes grouped in fault domain 1 is always 100% synced to the replica node and data nodes paired in fault domain 2. To activate CA, you must have at least one data node deployed, in addition to the primary node. If you have more than one data node, there must be an even number of data nodes including the primary node. For example, the cluster must have 2, 4, 6, 8, 10, 12, 14 or 16 nodes based on the appropriate sizing requirements. The data stored in the primary node in fault domain 1 is stored and replicated in the replica node in fault domain 2. The data stored in the data nodes in fault domain 1 is stored and replicated in the paired data nodes in fault domain 2. But in case the primary node fails, only the replica node can function as the replacement of the primary node.

- CA protects the analytics cluster against the loss of half the analytics nodes specific to one fault domain. You can stretch nodes across vSphere clusters in an attempt to isolate nodes or build failure zones.
- When CA is activated, the replica node can take over all functions that the primary node provides, in case of a primary node failure. The failover to the replica is automatic and requires only two to three minutes of VMware Aria Operations downtime to resume operations and restart data collection.

**NOTE**
In case of a primary node failure, the replica node becomes the primary node, and the cluster runs in degraded mode. To fix this, perform any one of the following actions.
- Correct the primary node failure manually.
- Return to CA mode by replacing the primary node. Replacement nodes do not repair the node failure, instead a new node assumes the primary node role.

- In the administration interface, after a CA replica node takes over and becomes the new primary node, you cannot remove the previous, offline primary node from the cluster. In addition, the previous node remains listed as a primary node. To refresh the display and activate the removal of the node, refresh the browser.
- When CA is activated, the cluster can survive the loss of half the data nodes, all in one fault domain, without losing any data. CA protects against the loss of only one fault domain at a time. Simultaneously losing data and primary/replica nodes, or two or more data nodes in both fault domains, is not supported.
- A CA activated cluster will be non-functional if you power off the primary node or the primary node replica while one of the fault domains is down.
- When CA is activated, it lowers the VMware Aria Operations capacity and processing by half, because CA creates a redundant copy of data throughout the cluster, and the replica backup of the primary node. Consider your potential use of CA when planning the number and size of your VMware Aria Operations cluster nodes. See Sizing the Cluster.
- When CA is activated, deploy analytics cluster nodes, in each fault domain, on separate hosts for redundancy and isolation. You can also use anti-affinity rules that keep nodes on specific hosts in the vSphere clusters.
- If you cannot keep the nodes separate in each fault domain, you can still activate CA. A host fault might cause the loss of the data nodes in the fault domain, and VMware Aria Operations can still be available in the other fault domain.
- If you cannot split the data nodes into different vSphere clusters, do not activate CA. A cluster failure can cause the loss of more than half of the data nodes, which is not supported, and all of vSphere might become unavailable.
- Without CA, you can keep nodes on the same host in the same vSphere. Without CA, the loss of even one node might make all of VMware Aria Operations unavailable.
- When you power off data nodes in both fault domains and change the network settings of the VMs, it affects the IP address of the data nodes. After this point, the CA cluster is no longer accessible and all the nodes status change to `"Waiting for analytics"`. Verify that you have used a static IP address.
- When you remove a node that has one or more vCenter adapters configured to collect data from a CA-activated cluster, one or more vCenter adapters associated with that node stops collecting. You must change the adapter configuration to pin them to another node before removing the node.
- The administration interface displays the resource cache count, which is created for active objects only, but the inventory displays all objects. When you remove a node from a CA-activated cluster allowing the vCenter adapters to collect data and rebalance each node, the inventory displays a different quantity of objects from that shown in the administration interface.

## Preparing for Installation

When you prepare for your installation, consider some of these best practices, cluster, sizing and scaling requirements.

## Requirements

You have to consider important requirements while creating nodes in a VMware Aria Operations.

### Using IPv6 with VMware Aria Operations

Requirements for IPv6
VMware Aria Operations supports both, Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6). All nodes in the cluster must follow the same protocol. For endpoint communications, you can use IPv4 or IPv6. If the environment

only supports the IPv6 protocol, the **Prefer IPv6** flag must be activated during the OVF deployment for each node. If you set the **Prefer IPv6** flag, then VMware Aria Operations uses IPv6 for all communications.

**Considerations While Using IPv6**

- If any nodes use DHCP, your DHCP server must be configured to support IPv6.
- IPv6 DHCP or Static configuration must have Global Scope.
- DHCP is only supported on data nodes. Primary nodes and replica nodes require static addresses.
- Your DNS server must be configured to support IPv6.
- When adding nodes to the cluster, enter the IPv6 address of the primary node.
- When registering a VMware vCenter instance within VMware Aria Operations, place square brackets around the IPv6 address of your VMware vCenter Server system if vCenter is also using IPv6.
  For example: `[2015:0db8:85a3:0042:1000:8a2e:0360:7334]`

  **NOTE**
  When VMware Aria Operations is using IPv6, vCenter Server might still have an IPv4 address. In that case, VMware Aria Operations does not need the square brackets.

- When registering a VMware vCenter® instance within VMware Aria Operations, place square brackets around the IPv6 address of your VMware vCenter Server® system if vCenter is also using IPv6.
  For example: `[2015:0db8:85a3:0042:1000:8a2e:0360:7334]`

  **NOTE**
  When VMware Aria Operations is using IPv6, vCenter Server might still have an IPv4 address. In that case, VMware Aria Operations does not need the square brackets.

### Cluster Requirements

When you create the cluster nodes that make up VMware Aria Operations, you have general requirements that you must meet.

### General VMware Aria Operations Cluster Node Requirements

General Cluster Node Requirements
You have to follow some general requirements to create a node on your environment.

**General Requirements**

- VMware Aria Operations version. All nodes must run the same VMware Aria Operations version.
  For example, do not add a version 6.1 data node to a cluster of VMware Aria Operations 6.2 nodes.
- Analytics Cluster Deployment Type. In the analytics cluster, all nodes must be the same kind of deployment.
- Witness Node Deployment Type. The witness node must be the same kind of deployment.
- Analytics Cluster Node Sizing. In the analytics cluster, CPU, memory, and disk size must be identical for all nodes. Primary, replica, and data nodes must be uniform in sizing.
- Witness Node Sizing. The witness node has only one size and may be of different sizes from the uniform analytics cluster node size
- Geographical Proximity. You may place analytics cluster nodes in different vSphere clusters, but the nodes must reside in the same geographical location.
  Different geographical locations are not supported.
- Witness Node Placement. You may place the witness node in a different vSphere cluster separate from the analytics nodes.

**NOTE**
A VMware Aria Operations cluster can have only one witness node.

- Virtual Machine Maintenance. When any node is a virtual machine, you may only update the virtual machine software by directly updating the VMware Aria Operations software.
  For example, going outside of VMware Aria Operations to access vSphere to update VMware Tools is not supported.
- Redundancy and Isolation. If you expect to activate HA, place analytics cluster nodes on separate hosts. See About High Availability .
- If you expect to activate CA, place analytics cluster nodes on separate hosts in fault domains, stretched across vSphere clusters. See About Continuous Availability.

**Requirements for Solutions**

Be aware that solutions might have requirements beyond those for VMware Aria Operations itself.

See your solution documentation, and verify any additional requirements before installing solutions. Note that the terms *solution*, *management pack*, *adapter*, and *plug-in* are used interchangeably.

### How VMware Aria Operations Uses Network Ports

How VMware Aria Operations Uses Network Ports
VMware Aria Operations uses network ports to communicate with a VMware vCenter Server system and VMware Aria Operations components.

In Linux deployments, you must manually verify or configure ports.

**IMPORTANT**
VMware Aria Operations does not support the customization of server ports.

Network Ports
Configure firewalls so that the following ports are open for bidirectional traffic.

**Table 18: Network Port Access Requirements for VMware Aria Operations**

| Port Number | Description |
|---|---|
| 22 (TCP) | Used for SSH access to the VMware Aria Operations cluster. |
| 80 (TCP) | Redirects to port 443. |
| 123 (UDP) | Used by VMware Aria Operations for Network Time Protocol (NTP) synchronization to the master node. |
| 443 (TCP) | Used to access the VMware Aria Operations product user interface and the VMware Aria Operations administrator interface. |
| 10443 (TCP) | Used by VMware Aria Operations to communicate with the vCenter Inventory service. |
| 3091–3094 (TCP) | When Horizon View (V4V) is installed, used to access data for VMware Aria Operations from V4V. |
| 5433 (TCP) | When high availability is enabled, used by the master and replica nodes to replicate the global database. |

*Table continued on next page*

*Continued from previous page*

| Port Number | Description |
|---|---|
| 6061 (TCP) | Used by clients to connect to the GemFire Locator to get connection information to servers in the distributed system. Also monitors server load to send clients to the least-loaded servers. |
| 7001 (TCP) | Used by Cassandra for secure inter-node cluster communication. |
| 9042 (TCP) | Used by Cassandra for secure client related communication amongst nodes. |
| 10000–10010 (TCP and UDP) | GemFire Server ephemeral port range used for unicast UDP messaging and for TCP failure detection in the peer-to-peer distributed system. |
| 20000–20010 (TCP and UDP) | GemFire Locator ephemeral port range used for unicast UDP messaging and for TCP failure detection in the peer-to-peer distributed system. |

Localhost Ports

Verify that your port configuration allows localhost access to the following ports. You may restrict off-host access to these ports if site policies are a concern.

**Table 19: Localhost Port Access Requirements for VMware Aria Operations**

| Port Number | Description |
|---|---|
| 1099 | GemFire Locator Java Management Extensions (JMX) Manager |
| 9004 | Analytics JMX Manager |
| 9008 | Cassandra database JMX Manager |
| 9160 | Cassandra Thrift client port |

### VMware Aria Operations Cluster Node Networking Requirements

Cluster Node Networking Requirements

When you create the cluster nodes that make up VMware Aria Operations, the associated setup within your network environment is critical to the inter-node communication and proper operation.

Networking Requirements

> **IMPORTANT**
> VMware Aria Operations analytics cluster nodes need frequent communication with one another. In general, your underlying vSphere architecture might create conditions where some vSphere actions affect that communication. Examples include, but are not limited to, vMotions, storage vMotions, HA events, and DRS events.

- The primary and replica nodes must use a static IP address, or fully qualified domain name (FQDN) with a static IP address.
  Data nodes can use dynamic host control protocol (DHCP).
- You can successfully reverse-DNS all nodes to their FQDN, currently the node hostname.
  Nodes deployed by OVF have their hostnames set to the retrieved FQDN by default.
- All nodes, must be bidirectionally routable by IP address or FQDN.

- Do not separate analytics cluster nodes with network address translation (NAT), load balancer, firewall, or a proxy that inhibits bidirectional communication by IP address or FQDN.
- Analytics cluster nodes must not have the same hostname.
- Place analytics cluster nodes within the same data center and connect them to the same local area network (LAN).
- Place analytics cluster nodes on same Layer 2 network and IP subnet.
  A stretched Layer 2 or routed Layer 3 network is not supported.
- Do not span the Layer 2 network across sites, which might create network partitions or network performance issues.
- With Continuous Availability activated, separate analytics cluster nodes into fault domains, stretched across vSphere clusters
- Packet Round Trip Time between the analytics cluster nodes must be 5 ms or lower.
- Network bandwidth between the analytics cluster nodes must be one gbps or higher.
- Do not distribute analytics cluster nodes over a wide area network (WAN).
  To collect data from a WAN, a remote or separate data center, or a different geographic location, use cloud proxy.
- Do not include an underscore in the hostname of any cluster node.
- Cloud proxies must have a proper DNS resolution to the VMware Aria Operations nodes when using short/long FQDN names. This is applicable to on-prem cloud proxy.

### VMware Aria Operations Cluster Node Best Practices

Cluster Node Best Practices
When you create the cluster nodes that make up VMware Aria Operations, additional best practices improve performance and reliability in VMware Aria Operations.

**Best Practices**

- Deploy VMware Aria Operations analytics cluster nodes in the same vSphere cluster in a single data center and add only one node at a time to a cluster allowing it to complete before adding another node.
- If you deploy analytics cluster nodes in a highly consolidated vSphere cluster, you might need resource reservations for optimal performance.
  Determine whether the virtual to physical CPU ratio is affecting performance by reviewing CPU ready time and co-stop.
- Deploy analytics cluster nodes on the same type of storage tier.
- To continue to meet analytics cluster node size and performance requirements, apply storage DRS anti-affinity rules so that nodes are on separate datastores.
- To prevent unintentional migration of nodes, set storage DRS to manual.
- To ensure balanced performance from analytics cluster nodes, use ESXi hosts with the same processor frequencies. Mixed frequencies and physical core counts might affect analytics cluster performance.
- To avoid a performance decrease, VMware Aria Operations analytics cluster nodes need guaranteed resources when running at scale. The VMware Aria Operations Knowledge Base includes sizing spreadsheets that calculate resources based on the number of objects and metrics that you expect to monitor, use of HA, and so on. When sizing, it is better to over-allocate than under-allocate resources.
  See Knowledge Base article 2093783.
- Because nodes might change roles, avoid machine names such as Primary, Data, Replica, and so on. Examples of changed roles might include making a data node into a replica for HA, or having a replica take over the primary node role.
- The NUMA placement is removed in the VMware Aria Operations 6.3 and later. Procedures related to NUMA settings from the OVA file follow:

**Table 20: NUMA Setting**

| Action | Description |
|---|---|
| Set the VMware Aria Operations cluster status to offline | 1. Shut down the VMware Aria Operations cluster.<br>2. Right-click the cluster and click **Edit Settings › Options › Advanced General**.<br>3. Click **Configuration Parameters**. In the vSphere Client, repeat these steps for each VM. |
| Remove the NUMA setting | 1. From the Configuration Parameters, remove the setting `numa.vcpu.preferHT` and click **OK**.<br>2. Click **OK**.<br>3. Repeat these steps for all the VMs in the VMware Aria Operations cluster.<br>4. Power on the cluster. |

**NOTE**
To ensure the availability of adequate resources and continued product performance, monitor VMware Aria Operations performance by checking its CPU usage, CPU ready and CPU contention time.

### Sizing and Scaling Requirements

The CPU, memory, and disk requirements that meet the needs of a particular environment depend on the number and type of objects in your environment and the data collected. This includes the number and type of adapters installed, the use of HA (High Availability) or CA (Continuous Availability), the duration of data retention, and the quantity of specific data points of interest.

VMware updates Knowledge Base article 2093783 with the most current information about sizing and scaling. The Knowledge Base article includes overall maximums and spreadsheet calculations that provide a recommendation based on the number of objects and metrics you expect to monitor.

## Installing VMware Aria Operations

VMware Aria Operations nodes are virtual appliance (vApp) based systems.

## Deployment of VMware Aria Operations

VMware Aria Operations consists of one or more nodes in a cluster. To create these nodes, you have to download and install the VMware Aria Operations suitable to your environment.

### Create a Node by Deploying an OVF

Create a Node by Deploying an OVF
VMware Aria Operations consists of one or more nodes, in a cluster. To create nodes, you use the vSphere client to download and deploy the VMware Aria Operations virtual machine, once for each cluster node.

- Verify that you have permissions to deploy OVF templates to the inventory.
- If the ESXi host is part of a cluster, activate DRS in the cluster. If an ESXi host belongs to a non-DRS cluster, all resource pool functions are deactivated.
- If this node is to be the primary node, reserve a static IP address for the virtual machine, and know the associated domain name, domain search path, domain name servers, default gateway, and network mask values.
  Plan to keep the IP address because it is difficult to change the address after installation.

- If this node is to be a data node that will become the HA/CA replica node, reserve a static IP address for the virtual machine, and store the associated domain name, domain search path, domain name servers, default gateway, and network mask values for later use.
  In addition, familiarize yourself with HA node placement as described in About High Availability and CA node allocation as described in About Continuous Availability .
- Plan your domain and machine naming so that the deployed virtual machine name begins and ends with an alphabet (a–z) or digit (0–9) characters, and will only contain alphabet, digit, or hyphen (-) characters. The underscore character (_) must not appear in the host name or anywhere in the fully qualified domain name (FQDN).
  Plan to keep the name because it is difficult to change the name after installation.

  For more information, review the host name specifications from the Internet Engineering Task Force. See www.ietf.org.
- Plan node placement and networking to meet the requirements described in General Cluster Node Requirements and Cluster Node Networking Requirements.
- If you expect the VMware Aria Operations cluster to use IPv6 addresses, review the IPv6 limitations described in Using IPv6 with .
- Download the VMware Aria Operations `.ova` file to a location that is accessible to the vSphere client.
- If you download the virtual machine and the file extension is `.tar`, change the file extension to `.ova`.
- Verify that you are connected to a vCenter system with a vSphere client, and log in to the vSphere client.
  Do not deploy VMware Aria Operations from an ESXi host. Deploy only from vCenter.

1. Select the vSphere**Deploy OVF Template** option.
2. Enter the path to the VMware Aria Operations `.ova` file.
3. Follow the prompts until you are asked to enter a name for the node.
4. Enter a node name. Examples might include `Ops1`, `Ops2Ops-A`, `Ops-B`.

   Do not include nonstandard characters such as underscores (_) in node names.

   Use a different name for each VMware Aria Operations node.

5. Follow the prompts until you are asked to select a configuration size.
6. Select the size configuration that you need. Your selection does not affect the disk size.

   Default disk space is allocated regardless of which size you select. If you need additional space to accommodate the expected data, add more disk after deploying the vApp, see Add Data Disk Space to a vApp Node.

7. Follow the prompts until you are asked to select the disk format.

| Option | Description |
|---|---|
| **Thick Provision Lazy Zeroed** | Creates a virtual disk in a default thick format. |
| **Thick Provision Eager Zeroed** | Creates a type of thick virtual disk that supports clustering features such as Fault Tolerance. Thick provisioned eager-zeroed format can improve performance depending on the underlying storage subsystem.<br><br>Select the thick provisioned eager-zero option when possible. |
| **Thin Provision** | Creates a disk in thin format. Use this format to save storage space. |

Snapshots can negatively affect the performance of a virtual machine and typically result in a 25–30 percent degradation for the VMware Aria Operations workload. Do not use snapshots.

8. Click **Next**.
9. From the drop-down menu, select a Destination Network, for example, **Network 1 = TEST**, and click **Next**.
10. Under Networking Properties, in case of a static IPv4 or static IPv6, specify the associated **Domain Name**, **Domain Search Path**, and **Domain Name Servers** values.
    a) Under the IPv4 configuration settings, specify the **IPv4 Type**. The primary node and replica node require a static IP. A data node can use DHCP or a static IP. From the drop down select one of the following.
       - **DHCP**: Select the DHCP type and leave all the fields blank.
       - **Static**: Select the Static type and enter the **IPv4 network IP Address**, **IPv4 default Gateway**, and the **IPv4 network Netmask values**.
       - **Disabled**: Disable IPv4.

    b) Under the IPv6 configuration settings, click the **Prefer IPv6** checkbox to use IPv6 for cluster creation and internode communication. Specify the **IPv6 Type**. The primary node and replica node require a static IP. A data node can use DHCP or a static IP. From the drop down select one of the following.
       - **DHCP/Slaac**: Select DHCP/Slaac and leave all the fields blank.
       - **Static**: Select Static and enter the **IPv6 network IP Address**, **IPv6 default Gateway**, and the **IPv6 network Netmask** values.
       - **Disabled**: Disable IPv6.
11. In the Timezone Setting, leave the default of UTC or select a time zone.

    The preferred approach is to standardize on UTC. Alternatively, configure all nodes to the same time zone.

    **NOTE**
    You cannot configure nodes to different time zones.

12. If you want to deploy a FIPS activated VMware Aria Operations setup, in the FIPS setting, select the **Activate FIPS Mode** check box.
13. Click **Next**.
14. Review the settings and click **Finish**.
15. If you are creating a multiple-node VMware Aria Operations cluster, repeat through all the steps to deploy each node.

Use a Web browser client to configure a newly added node as the VMware Aria Operations primary node, a data node, or a high availability primary replica node. The primary node is required first.

**NOTE**
For security, do not access VMware Aria Operations from untrusted or unpatched clients, or from clients using browser extensions.

## Installation Types

After you have installed VMware Aria Operations product, you can either perform a new installation, an express installation, or expand an existing installation.

- Express Installation
- New installation
- Expand Installation

**Figure 2: Getting Started Setup**

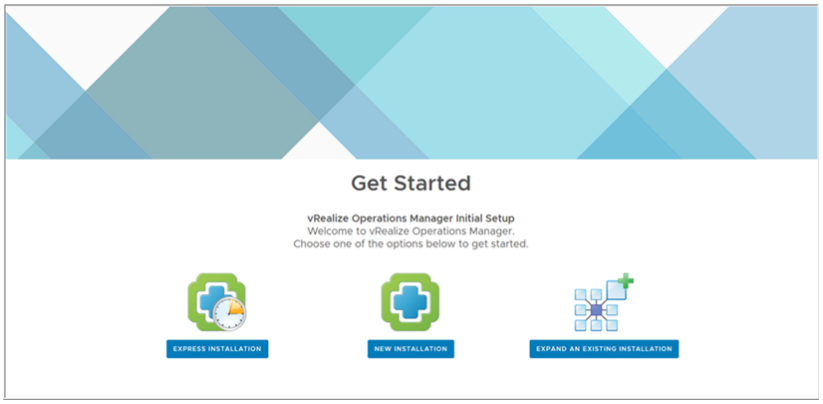


### Installing VMware Aria Operations for a New User

After you install VMware Aria Operations using an OVF or an installer, you are notified to the main product UI page. You can create a single node or multiple nodes depending on your environment.

### Introduction to a New Installation

You can perform a new installation as a first-time user and create a single node to handle both administration and data handling.

**Figure 3: New Installation from the Setup page**


### Perform a New Installation on the VMware Aria Operations Product UI

You can create a single node and configure it as a primary node or create a data node in a cluster to handle additional data. All VMware Aria Operations installations require a primary node. With a single node cluster, administration and data functions are on the same primary node. A multiple-node VMware Aria Operations cluster contains one primary node and one or more nodes for handling additional data.

- Create a node by deploying the VMware Aria Operations vApp.
- Create a node by running the VMware Aria Operations Enterprise installer for Linux.
- Alternatively, create a node by running the VMware Aria Operations Enterprise installer for Linux.
- After it is deployed, note the fully qualified domain name (FQDN) or IP address of the node.
- If you plan to use a custom authentication certificate, verify that your certificate file meets the requirements for VMware Aria Operations.

1. Navigate to the name or IP address of the node that will be the primary node of VMware Aria Operations.

   The setup wizard appears, and you do not need to log in to VMware Aria Operations.

2. Click **New Installation**.
3. Click **Next**.
4. Enter and confirm a password for the admin user account, and click **Next**.

   Passwords require a minimum of eight characters, one uppercase letter, one lowercase letter, one digit, and one special character.

   The user account name is admin by default and cannot be changed.

5. Select whether to use the certificate included with VMware Aria Operations or to install one of your own.
   a) To use your own certificate, click **Browse**, locate the certificate file, and click **Open** to load the file in the Certificate Information text box.
   b) Review the information detected from your certificate to verify that it meets the requirements for VMware Aria Operations.
6. Click **Next**.
7. Enter a name for the primary node.
   For example: `Ops-Primary`
8. Enter the URL or IP address for the Network Time Protocol (NTP) server with which the cluster synchronizes.
   For example: `nist.time.gov`
9. Click **Add**.

   Leave the NTP blank to have VMware Aria Operations manage its own synchronization by having all nodes synchronize with the primary node and replica node.

10. Click **Next**.
11. Configure the VMware Aria Operations availability. To install VMware Aria Operations with availability, activate the **Availability Mode** and select High Availability or Continuous Availability. To continue your installation on full capacity, click **Next** .

    **NOTE**
    You can activate High Availability or Continuous Availability after installation from the administrator interface.

12. Click the Add icon to add a node.
    a) Enter the **Node Name** and **Node Address**.
    b) Select the **Current Cluster Role**.

    **NOTE**
    This step is optional if you use the default configuration. If you select High Availability for this cluster option, you can select a node from the added list of nodes to be the replica node. Although, only one node from the list can be selected as a replica node. For more information on High Availability, see Adding High Availability to . If you select Continuous Availability for this cluster, add at least one witness node and an even number of data nodes including the primary node and divide them across two fault domains. For more information, see Adding Continuous Availability.

13. Click **Next**, and click **Finish**.
    The administration interface appears, and it takes a moment for VMware Aria Operations to finish adding the primary node.

You have created a primary node to which you can add more nodes.

After creating the primary node, you have the following options.

- Create and add data nodes to the unstarted cluster.
- Click **Start VMware Aria Operations** to start the single-node cluster, and log in to finish configuring the product.
  The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.

  If you require a new agent to collect data, you must deploy a cloud proxy. For more information on how to deploy a cloud proxy, see Installing .

### About the VMware Aria Operations Primary Node

About the Primary Node
The primary node is the required, initial node in your VMware Aria Operations cluster.

The primary node performs administration for the cluster and must be online before you configure any new nodes. In addition, the primary node must be online before other nodes are brought online. If the primary node and replica node go offline together, bring them back online separately. Bring the primary node online first, and then bring the replica node online.

### Advantages of a New Installation

You can use the new installation to create a primary node during the first installation of VMware Aria Operations. With the primary node in place, you can then start adding more nodes to form a cluster and then define an environment for your organization.

In a single-node clusters, administration and data is on the same primary node. A multiple-node cluster includes one primary node and one or more data nodes. In addition, there might be one replica node used for high availability. For continuous availability, you need a witness node and an even number of data nodes including the primary node. For more information on creating a primary node, see About the Primary Node.

### Installing VMware Aria Operations as an Administrator

As an administrator, you can install several instances of VMware Aria Operations build in your VM environment.

### Introduction to Express Installation

Express installation is one possible way to create primary nodes, add data nodes, form clusters, and test your connection status. You can use express installation to save time and speed up the process of installation when compared to a new installation. Do not to use this feature unless the user is an administrator.

**Figure 4: Express Installation from the Setup screen**


### Perform an Express Installation on the VMware Aria Operations product UI

Use express installation on the VMware Aria Operations cluster to create a primary node. Select express installation option when installing for the first time.

Verify that you have a static IP address created from an OVF file.

1. Navigate to the name or IP address of the node that will be the primary node of VMware Aria Operations.

   The setup wizard appears, and you do not need to log in to VMware Aria Operations.
2. Click **Express Installation**.
3. Click **Next**.
4. Enter and confirm a password for the admin user account, and click **Next**.

   Passwords require a minimum of 8 characters, one uppercase letter, one lowercase letter, one digit, and one special character.

   The user account name is admin by default and cannot be changed.
5. Click **Next**.
6. Click **Finish**.

You have created a primary node to which you can add more nodes.

### Advantages of an Express Installation

Express installation saves time when compared to a new installation to create a new primary node. The express installation uses the default certificates, which differ from one organization to another. This feature is mainly used by the developers or the administrators.

### Expand an Existing Installation of VMware Aria Operations

Use this option to add a node to an existing VMware Aria Operations cluster. You can use this option if you have already configured a primary node and you want to increase the capacity by adding more nodes to your cluster.

## Introduction to Expand an Existing Installation

You can deploy and configure additional nodes so that VMware Aria Operations can support larger environments. A primary node always requires an additional node for a cluster to monitor your environment. With expanding your installation, you can add more than one node to your cluster.

### Adding Data Nodes

Data nodes are the additional cluster nodes that allow you to scale out VMware Aria Operations to monitor larger environments.

You can dynamically scale out VMware Aria Operations by adding data nodes without stopping the VMware Aria Operations cluster. When you scale out the cluster by 25% or more, you should restart the cluster to allow VMware Aria Operations to update its storage size, and you might notice a decrease in performance until you restart. A maintenance interval provides a good opportunity to restart the VMware Aria Operations cluster.

In addition, the product administration options include an option to re-balance the cluster, which can be done without restarting. Rebalancing adjusts the VMware Aria Operations workload across the cluster nodes.

**Figure 5: Expand an existing installation from the Setup screen**

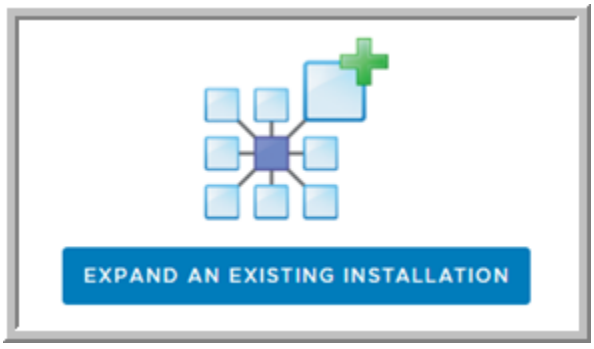


**NOTE**
Do not shut down online cluster nodes externally or by using any means other than the VMware Aria Operations interface. Shut down a node externally only after taking it offline in the VMware Aria Operations interface.

### Expand an Existing Installation to Add a Data Node

Expand an existing installation to add a data node
Larger environments with multiple-node VMware Aria Operations clusters contain one primary node and one or more data nodes for additional data collection, storage, processing, and analysis.

- Create nodes by deploying the VMware Aria Operations vApp.
- Create nodes by running the VMware Aria Operations Enterprise installer for Linux.
- Alternatively, create nodes by running the VMware Aria Operations Enterprise installer for Linux.
- Create and configure the primary node.
- Note the fully qualified domain name (FQDN) or IP address of the primary node.

1. In a Web browser, navigate to the name or IP address of the node that will become the data node.

   The setup wizard appears, and you do not need to log in to VMware Aria Operations.

2. Click **Expand an Existing Installation**.
3. Click **Next**.

4. Enter a name for the node (for example, `Data-1`).
5. From the Node Type drop-down, select **Data**.
6. Enter the FQDN or IP address of the primary node and click **Validate**.
7. Select **Accept this certificate** and click **Next**.

   If necessary, locate the certificate on the primary node and verify the thumbprint.

8. Verify the VMware Aria Operations administrator username of admin.
9. Enter the VMware Aria Operations administrator password.

   Alternatively, instead of a password, type a pass-phrase that you were given by your VMware Aria Operations administrator.

10. Click **Next**, and click **Finish**.
    The administration interface appears, and it takes a moment for VMware Aria Operations to finish adding the data node.

After creating a data node, you have the following options.

- New, unstarted clusters:
  - Create and add more data nodes.
  - Create a high availability primary replica node.
  - In a Web browser, navigate to the primary node administration interface at `https://primary-node-name-or-ip-address/admin`. Verify that all the nodes are listed under the **Nodes in the VMware Aria Operations Cluster**. Then, click **Start VMware Aria Operations** to start the cluster and to finish configuring the product. The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.
- Established, running clusters:
  - Create and add more data nodes.
  - Create a high availability primary replica node, which requires a cluster restart.

#### Advantages of an Expanding an Installation

A data node shares the load of performing VMware Aria Operations analysis and it can also have an adapter installed to perform collection and data storage from the environment. You must have a primary node before you add data nodes to form a cluster.

## Using VMware Aria Operations on-premises to Monitor VMware Cloud

You can use your VMware Aria Operations on-premises to manage and monitor your cloud infrastructure on VMware Clouds by simply adding a dedicated cloud account for VMware Cloud or by adding a vCenter Server account in VMware Aria Operations. You can extend the current set of monitoring, troubleshooting, optimization, and remediation processes of VMware Aria Operations on to VMware Cloud. It also provides you with a hybrid view of your environment.

**Prerequisites**

- A VPN or a direct connection to set up the bidirectional access between the nodes and cloud proxies of VMware Aria Operations on-premises and VMware Cloud.
- Scale the existing VMware Aria Operations cluster before adding the new VMware Cloud SDDC sites. https://vropssizer.vmware.com/.

## VMware Cloud on AWS

The following diagram shows VMware Aria Operations on-premises collecting data from VMware Cloud on AWS with cloud proxy. Configure VMware Aria Operations to monitor VMware Cloud on AWS using the steps described in the topic 'Configuring VMware Cloud on AWS in VMware Aria Operations' in the *VMware Aria Operations Configuration Guide*.

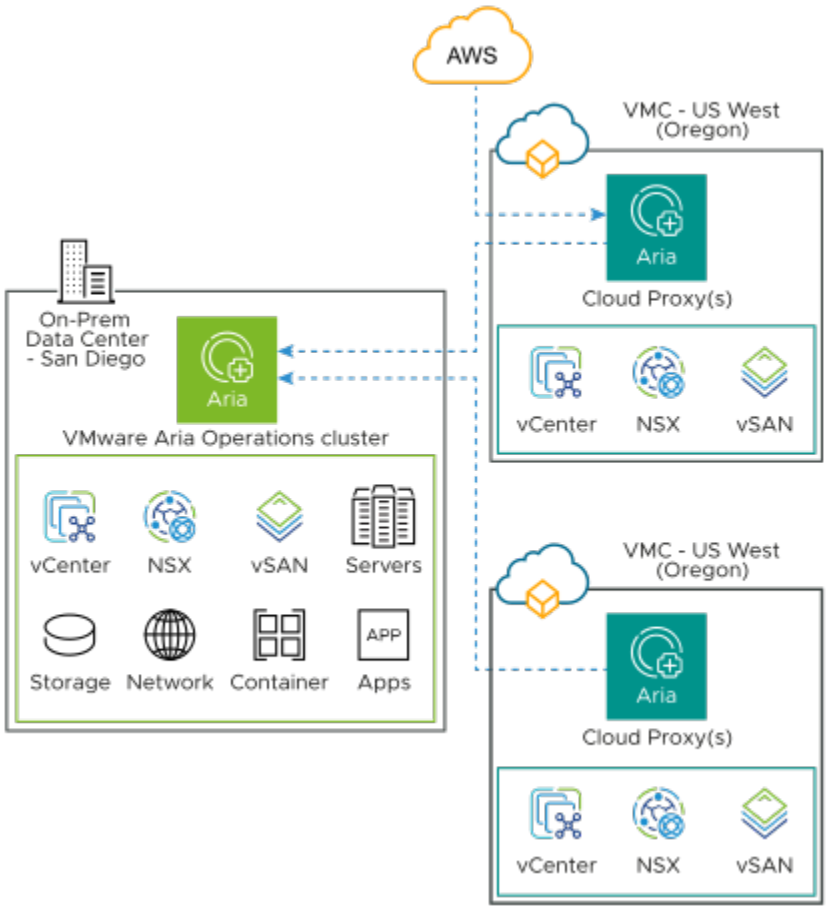


## Azure VMware Solution

The following diagram shows VMware Aria Operations on-premises collecting data from Azure VMware Solution with cloud proxies. Configure VMware Aria Operations to monitor Azure VMware Solution using the steps described in the topic Configuring an Azure VMware Solution Instance in VMware Aria Operations.

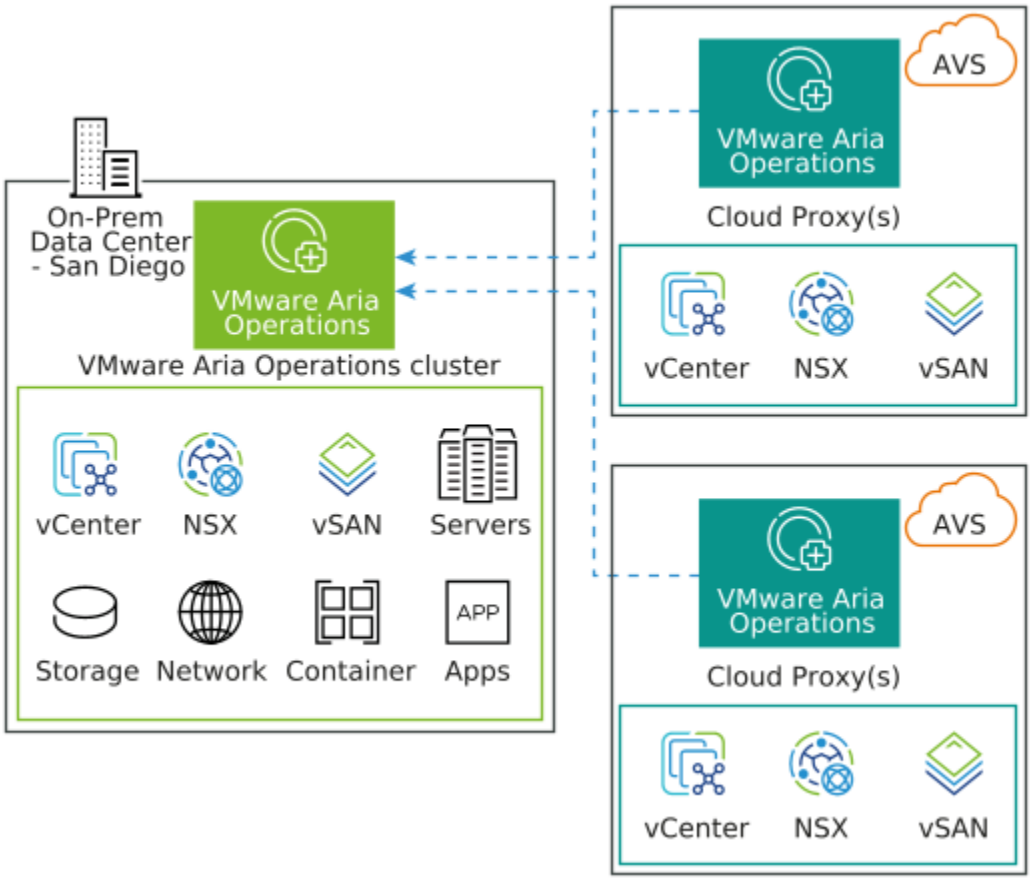


**Google Cloud VMware Engine**

The following diagram shows VMware Aria Operations on-premises collecting data from Google Cloud VMware Engine with cloud proxies. Configure VMware Aria Operations to monitor Google Cloud VMware Engine using the steps described in the topic Configure a Google Cloud VMware Engine Instance in VMware Aria Operation.

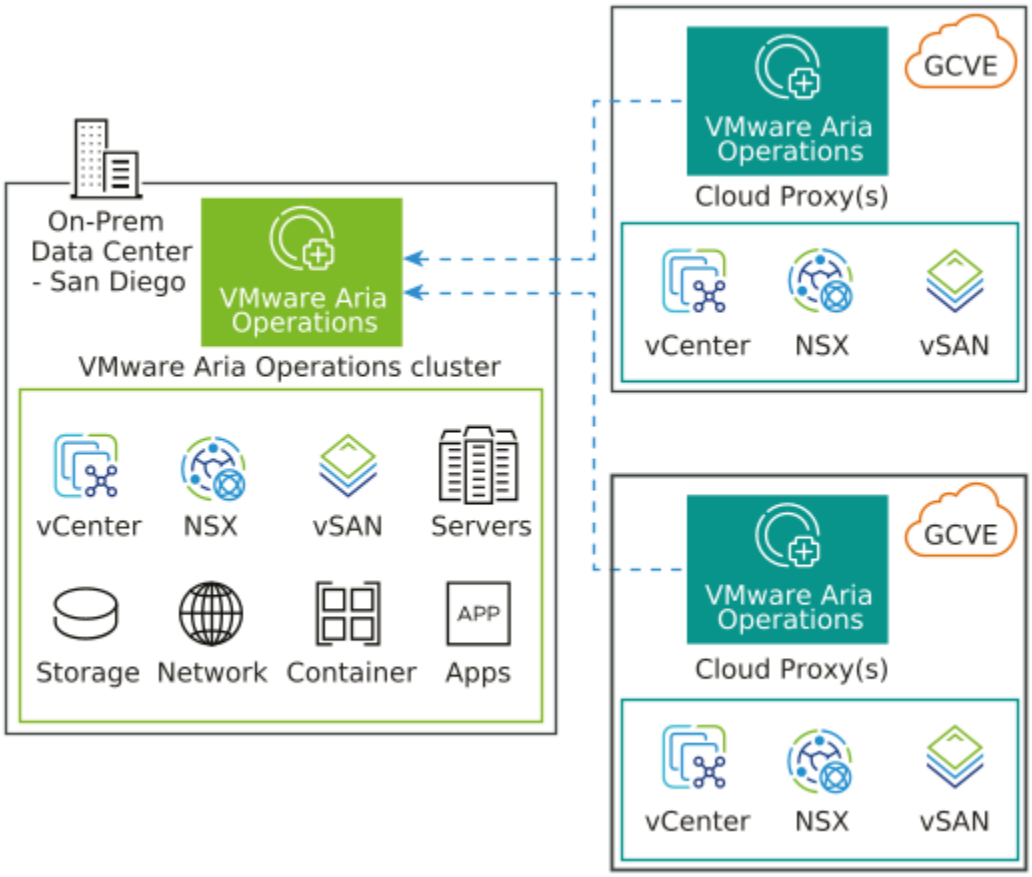


## Resize your Cluster by Adding Nodes

You can deploy and configure additional nodes so that VMware Aria Operations can support larger environments.

**Figure 6: Workflow - Resize your cluster**

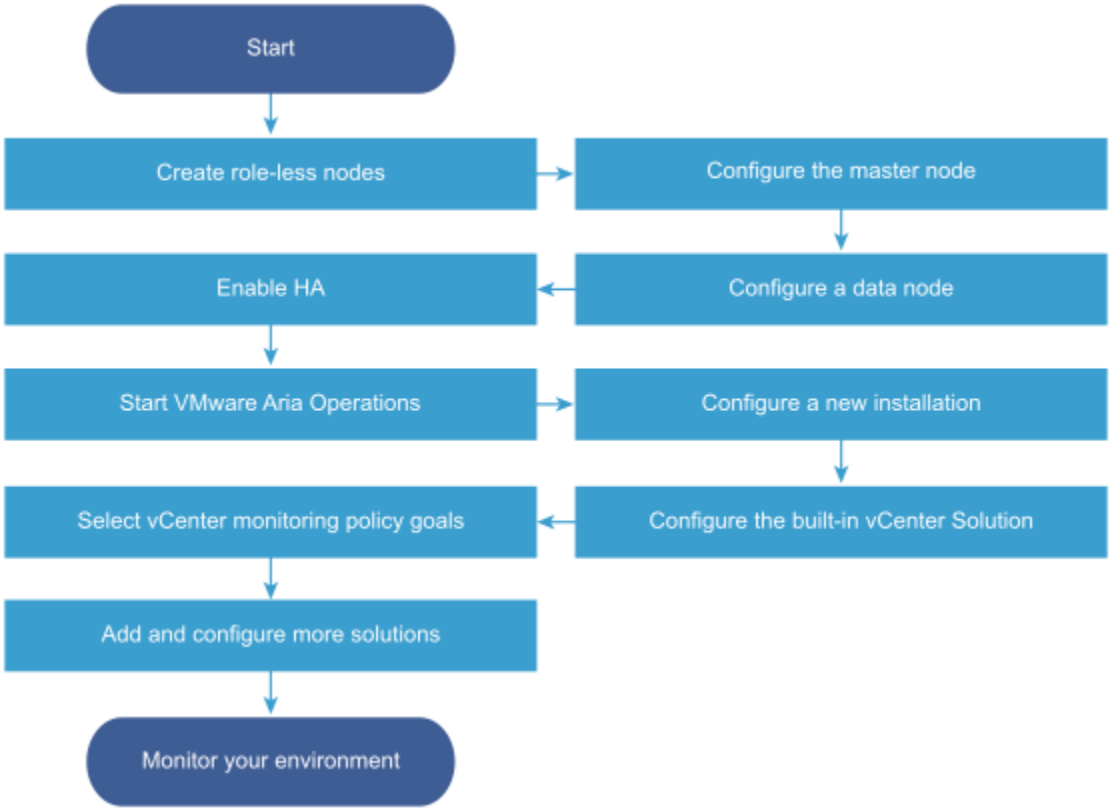


## Adding High Availability to VMware Aria Operations

Adding High Availability
You can dedicate one VMware Aria Operations cluster node to serve as a replica node for the VMware Aria Operations primary node.

### Run the Setup Wizard to Add a Primary Replica Node

Run the Setup Wizard to Add a Primary Replica Node
To activate high availability (HA) for a VMware Aria Operations cluster, specify one of the data nodes to become a replica of the primary node.

- Create nodes by deploying the VMware Aria Operations vApp.
- Create nodes by running the VMware Aria Operations Enterprise installer for Linux.
- Alternatively, create nodes by running the VMware Aria Operations Enterprise installer for Linux.
- Create and configure the primary node.
- Create and configure a data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the primary node.

> **NOTE**
> If the cluster is running, activating HA restarts the cluster.

You can add HA to the VMware Aria Operations cluster at installation time or after VMware Aria Operations is up and running. Adding HA at installation is less intrusive because the cluster has not yet started. For more information, see the

"High Availability Considerations" topic in the *Reference Architecture* guide and the "High Availability (HA)" topic in the *VMware Aria Operations Best Practices* guide.

1. In a Web browser, navigate to the primary node administration interface.
   `https://primary-node-name-or-ip-address/admin`
2. Enter the VMware Aria Operations administrator user name of `admin`.
3. Enter the VMware Aria Operations administrator password and click **Log In**.
4. Under High Availability, click **Activate**.
5. Select a data node to serve as the replica for the primary node.
6. Select the **Activate High Availability for this cluster** option, and click **OK**.
   If the cluster was online, the administration interface displays progress as VMware Aria Operations configures, synchronizes, and rebalances the cluster for HA.
7. If the primary node and replica node go offline, and the primary remains offline for any reason while the replica goes online, the replica node does not take over the primary role, take the entire cluster offline, including data nodes and log in to the replica node command-line console as a root.
8. Open `$ALIVE_BASE/persistence/persistence.properties` in a text editor.
9. Locate and set the following properties:
   `db.role=PRIMARY`

   `db.driver=/data/vcops/xdb/vcops.bootstrap`
10. Save and close *persistence.properties*.
11. In the administration interface, bring the replica node online, and verify that it becomes the primary node and bring the remaining cluster nodes online.

After creating a primary replica node, you have the following options.

- New, unstarted clusters:
  - Create and add data nodes.
  - Click **Start VMware Aria Operations** to start the cluster, and log in to finish configuring the product.
    The cluster might take from 10 to 30 minutes to start, depending on the size of your cluster and nodes. Do not make changes or perform any actions on cluster nodes while the cluster is starting.
- Established, running clusters:
  - Create and add data nodes.
- To deactivate high availability, click **Deactivate HA** and follow the steps in the admin UI. The deativation timeline depends on the cluster size and the data retention period. Please wait for the deactivation to complete.

## Adding Continuous Availability

Continuous availability prevents data loss in the event of one or more node failures. This mode requires one witness node, one primary node, and one data node divided across two fault domains. The witness node lies outside the fault domains. By default, the primary node is assigned to **Fault Domain 1**. The data node becomes the replica node and is assigned to **Fault Domain 2**. The primary node and the replica node create a pair. The number of data nodes including the primary node should always be an even number not exceeding 16. Each data node added to **Fault Domain 1** must have a pair in **Fault Domain 2** to preserve and replicate data that is added to its peer.

### Activate Continuous Availability in VMware Aria Operations

You can activate continuous availability (CA) for VMware Aria Operations to protect your data if there is one or more node failures.

- Create nodes by deploying the VMware Aria Operations vApp.

- Create nodes by running the VMware Aria Operations Enterprise installer for Linux.
- Alternatively, create nodes by running the VMware Aria Operations Enterprise installer for Linux.
- Create and configure the primary node.
- Create and configure the witness node.

  **NOTE**
  VMware Aria Operations can have only one witness node in its cluster. While deploying an OVA file, you can select the recommended CPU/RAM configuration for the witness node. For sizing information, see KB Article 2093783.

- Create and configure one data node with a static IP address.
- Note the fully qualified domain name (FQDN) or IP address of the primary node.

**NOTE**
CA can be activated only on a standard cluster. If the cluster is running, activating CA restarts the cluster.

You can activate CA in the VMware Aria Operations cluster at the installation time or after VMware Aria Operations is up and running. Adding CA at installation is less intrusive because the cluster has not yet started. For more information, see the "Continuous Availability Considerations" and "Continuous Availability FAQ" topics in the *Reference Architecture* guide and the "Continuous Availability (CA)" topic in the *VMware Aria Operations Best Practices* guide.

1. In a Web browser, navigate to the primary node administration interface.

   `https://primary-node-name-or-ip-address/admin`

2. Enter the VMware Aria Operations administrator user name of `admin`.
3. Enter the VMware Aria Operations administrator password and click **Log In**.
4. Under Continuous Availability, click **Activate CA**.
   The Continuous Availability wizard opens. The Witness node exists outside the fault domains. The primary node is already assigned to **Fault Domain 1**.

   **NOTE**
   You can enter names for each Fault Domain during installation. You can also edit the fault domain names after activating continous availability.

5. To create a pair with the primary node, drag the data nodes to **Fault Domain 2**.

   **NOTE**
   You can add a maximum of 16 data nodes including the primary node and divide them between the fault domains to create eight pairs.

6. Click **Ok**.

- To deactivate continuous availability, click **Deactivate CA** and follow the steps in the admin UI. The deativation timeline depends on the cluster size and the data retention period. Please wait for the deactivation to complete.

## VMware Aria Operations Cluster and Node Maintenance

Cluster and Node Maintenance
You perform cluster and node maintenance procedures to help your VMware Aria Operations perform more efficiently cluster and node maintenance involves activities such as changing the online or offline state of the cluster, fault domains, or individual nodes, activating or deactivating high availability (HA) or continuous availability (CA), reviewing statistics related to the installed adapters, and rebalancing the workload for a better performance.

You perform most VMware Aria Operations cluster and node maintenance using the Cluster Management page in the product interface, or the Cluster Status and Troubleshooting page in the administration interface. The administration interface provides more options than the product interface.

**Table 21: Cluster and Node Maintenance Procedures**

| Procedure | Interface | Description |
|---|---|---|
| Change Cluster Status | Administration/Product | You can change the status of a node to online or offline.<br>In a high availability (HA) cluster, taking the primary or replica offline causes VMware Aria Operations to run from the remaining node and for HA status to be degraded.<br><br>In continuous availability (CA) cluster, taking the primary or replica offline causes VMware Aria Operations to run in a degraded status.<br><br>**NOTE**<br>You cannot convert a High Availability (HA) activated cluster to a Continuous Availability cluster and vice versa. You must first deactivate the cluster availability, so that the cluster becomes a standard cluster and then activate HA or CA as required.<br><br>Any manual or system action that restarts the cluster brings all VMware Aria Operations nodes online, including any nodes that you had taken offline. |
| Activate or Deactivate High Availability | Administration | Activating high availability requires the cluster to have at least one data node, with all nodes online or all offline.<br><br>To activate high availability, see Adding High Availability to .<br><br>Deactivating high availability restarts the VMware Aria Operations cluster.<br><br>After you deactivate high availability, the replica node in VMware Aria Operations converts back to a data node and restarts the cluster. |
| Activate or Deactivate Continuous Availability | Administration | Activating continuous availability requires the cluster to have at least one witness node, and at least two data node, with all nodes online or all offline.<br><br>To activate continuous availability, see Adding Continuous Availability. |

*Table continued on next page*

*Continued from previous page*

| Procedure | Interface | Description |
|---|---|---|
| | | Deactivating continuous availability restarts the VMware Aria Operations cluster.<br><br>When you deactivate continuous availability, you can choose to keep all your nodes or cut out one of the fault domains.<br>• Click **Simply Deactivate with keeping all nodes** to keep all your nodes when you deactivate continuous availability.<br><br>**NOTE**<br>You cannot deactivate continuous availability if one of your nodes is faulty. If you want to keep all your nodes, you must fix or replace the faulty node before you proceed.<br><br>• Click **Cut-Out one Fault Domain** and then select the fault domain you want to keep. The other fault domain and the witness node are deleted.<br><br>After you deactivate continuous availability, the replica node in VMware Aria Operations converts back to a data node and restarts the cluster. |
| Add Nodes | Administration | You can add one or more nodes for your cluster.<br><br>In a FIPS activated environment, new nodes must be FIPS compliant. In a FIPS deactivated environment, new nodes must be FIPS deactivated.<br><br>Activating continuous availability requires one witness node, and an even number of data nodes including the primary node. For example, the cluster must have 2, 4, 6, 8, 10, 12, 14 or 16 nodes. |
| Replace Nodes | Administration | You can add nodes and replace them with a downed or non-functional node in a cluster. |
| Generate Passphrase | Administration | You can generate a passphrase to use instead of the administrator credentials to add a node to this cluster.<br>The passphrase is only valid for a single use. |
| Remove a Node | Administration | When you remove a node, you lose data that the node had collected unless you are running in high availability (HA) mode. HA protects against the removal or loss of one node.<br>You must not re-add nodes to VMware Aria Operations that you already removed. If your environment requires more nodes, add new nodes instead. |

*Table continued on next page*

*Continued from previous page*

| Procedure | Interface | Description |
|---|---|---|
| | | When you perform maintenance and migration procedures, you should take the node offline, not remove the node. |
| Configure NTP | Product | The nodes in VMware Aria Operations cluster synchronize with each other by standardizing on the primary node time or by synchronizing with an external Network Time Protocol (NTP) source. |
| Rebalance the Cluster | Product | You can rebalance adapter, disk, memory, or network load across VMware Aria Operations cluster nodes to increase the efficiency of your environment. |

## VMware Aria Operations Post-Installation Considerations

Post-Installation Considerations
After you install VMware Aria Operations, there are post-installation tasks that might need your attention.

## About Logging In to VMware Aria Operations

About Logging In
Logging in to VMware Aria Operations requires that you point a Web browser to the fully qualified domain name (FQDN) or IP address of a node in the VMware Aria Operations cluster.

When you log in to VMware Aria Operations, there are a few things to keep in mind.

- After initial configuration, the product interface URL is:
  https://*node-FQDN-or-IP-address*
- Before initial configuration, the product URL opens the administration interface instead.
- After initial configuration, the administration interface URL is:
  https://*node-FQDN-or-IP-address*/admin
- The administrator account name is admin. The account name cannot be changed.
- The admin account is different from the root account used to log in to the console, and does not need to have the same password.
- When logged in to the administration interface, avoid taking the node that you are logged into offline and shutting it down. Otherwise, the interface closes.
- The number of simultaneous login sessions before a performance decrease depends on factors such as the number of nodes in the analytics cluster, the size of those nodes, and the load that each user session expects to put on the system. Heavy users might engage in significant administrative activity, multiple simultaneous dashboards, cluster management tasks, and so on. Light users are more common and often require only one or two dashboards.
  The sizing spreadsheet for your version of VMware Aria Operations contains further detail about simultaneous login support. See Knowledge Base article 2093783.
- You cannot log in to a VMware Aria Operations interface with user accounts that are internal to VMware Aria Operations, such as the maintenance Admin account.
- For supported Web browsers, see the VMware Aria Operations Release Notes for your version.

## Logging In to VMware Aria Operations

Logging In
After installing a new instance of VMware Aria Operations, you must log in and complete a one-time process to license the product and configure solutions for the kinds of objects that you want to monitor.

- Create the new cluster of VMware Aria Operations nodes.
- Verify that the cluster has enough capacity to monitor your environment. See Sizing the Cluster.

1. In a Web browser, navigate to the IP address or fully qualified domain name of the primary node.
2. Enter the username `admin` and the password that you defined when you configured the primary node and click **Login**.
   The administration interface appears when you login for the first time.
3. To start the cluster, click **Start VMware Aria Operations** and then click **Yes**.

   The cluster might take from 10 to 30 minutes to start, depending on your environment. Do not make changes or perform any actions on the cluster nodes while the cluster is starting.

4. When the cluster starts and the product login page appears, enter the admin username and password again, and click **Login**.
   A one-time licensing wizard appears.
5. Click **Next**, Read and accept the End User License Agreement, and then click **Next**.
6. Enter your product key, or select the option to run VMware Aria Operations in evaluation mode.

   Your level of product license determines what solutions you may install to monitor and manage objects.
   - Standard. vCenter only
   - Advanced. vCenter plus other infrastructure solutions
   - Enterprise. All solutions

   VMware Aria Operations does not license managed objects in the same way that vSphere does, so there is no object count when you license the product.

   **NOTE**
   When you transition to the Standard edition, you no longer have the Advanced and Enterprise features. After the transition, delete any content that you created in the other versions to ensure that you comply with EULA and verify the license key which supports the Advanced and Enterprise features.

7. If you entered a product key, click **Validate License Key**, and then click **Next**
8. The Customer Experience Improvement Program (CEIP) sends usage statistics to VMware.

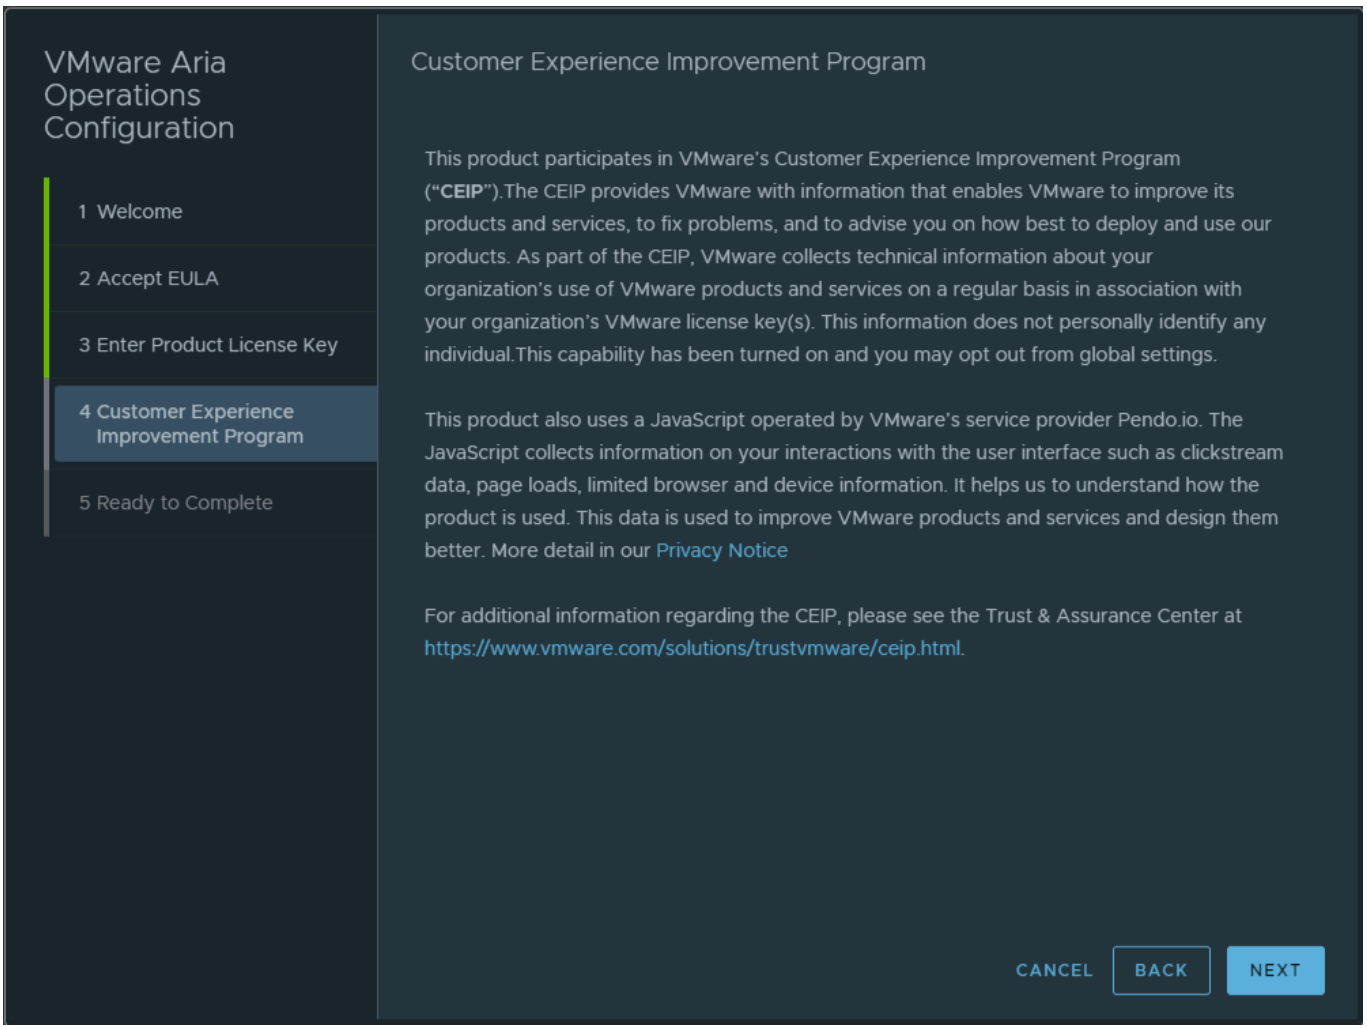


If you do not want to participate in CEIP, you can deactivate this setting from the Global Settings page. For more information, see 'List of Global Settings' in *Configuring VMware Aria Operations Guide*.

9. Click **Next** and then click **Finish**.
   The one-time wizard finishes, and the VMware Aria Operations interface appears.

- Use the VMware Aria Operations interface to configure the solutions that are included with the product.
- Use the VMware Aria Operations interface to add more solutions.
- Use the VMware Aria Operations interface to add monitoring policies.

## After You Log In

After you log in to VMware Aria Operations from a web browser, you see the newly designed Launchpad in the Home page. You can set any dashboard to appear in the home page, beside the Launchpad. Click the **Actions** menu on a dashboard that you want to set as the landing page and select **Set as Home landing page**. To remove the dashboard as the home landing page, click the **Actions** menu on the relevant dashboard and select **Reset from Home landing page**.

## Home page When You First Log In

If you have logged in using an administrative account, you must set the currency in the **Global Settings** page. From the left menu, click **Administration**, and then click the **Global Settings** tile. You can do so from the message that you see in the Launchpad when you log in for the first time. Optionally, you can close the message. Once you set a currency, you cannot change it. As an administrator, you must also first set up a cloud account or configure an adapter before you can start using VMware Aria Operations.

The following message in the Launchpad indicates that VMware Aria Operations is in the evaluation mode or grace period: `You are in an evaluation mode or in a grace period. Add a VMware Cloud Foundation or VMware vSphere Foundation license to a connected vCenter instance or use the legacy license model.` For more information on licensing, see the "About VMware Aria Operations Licenses" topic in the *VMware Aria Operations Configuration Guide*.

After logging in you land on the Launchpad and if you see a message like, "`VMware Aria Operations internal certificates will expire on dd/mm/yyyy. Please install a new certificate before the expiry date. For details, see KB 71018`", you must upgrade your internal certificates for VMware Aria Operations using the certificate renewal PAK file from the VMware Aria Operations Administrator interface. For more information, see the KB article 71018.

## Launchpad After Cloud Accounts Are Configured

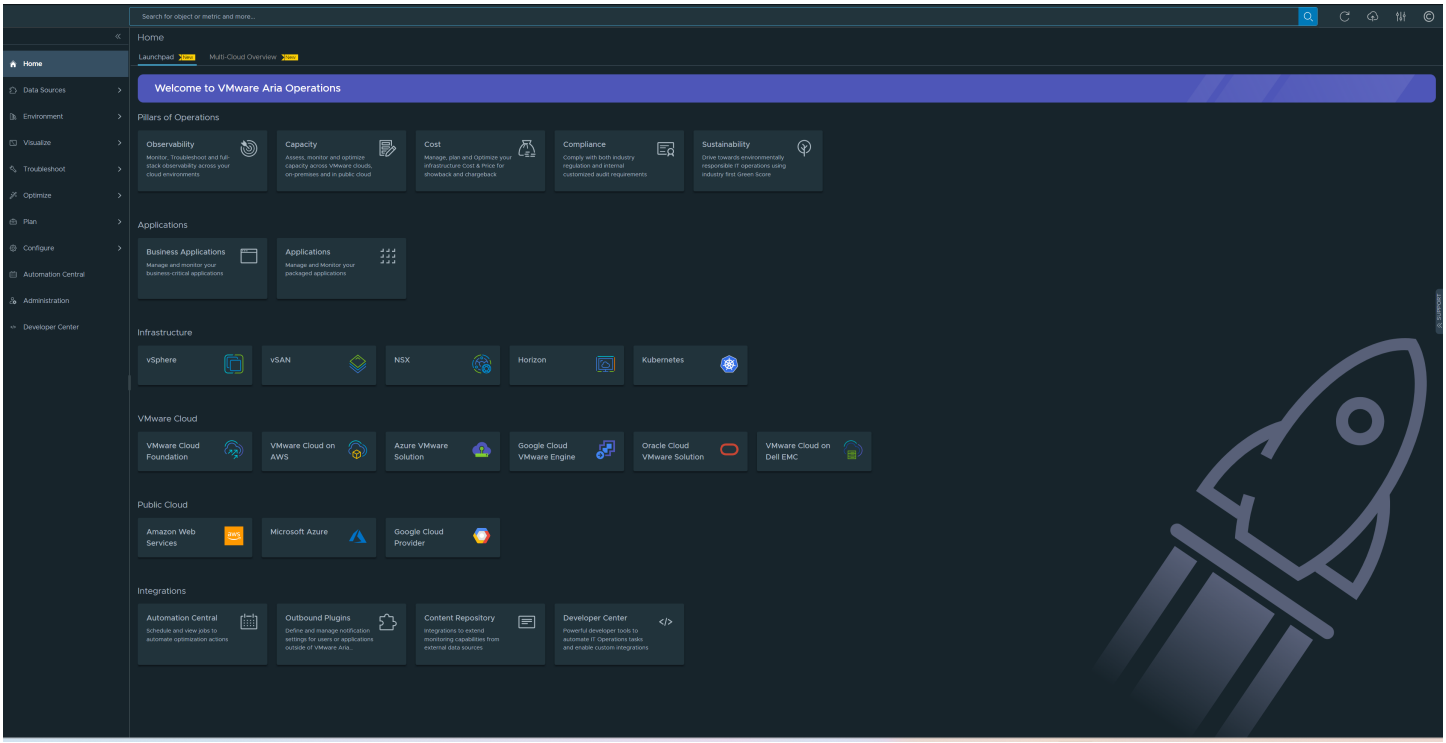


Beside the **Launchpad** tab, you see the **Multi-Cloud Overview** tab. To know more about the Multi-Cloud Overview, see the topic, Monitoring Multiple Cloud Accounts in VMware Aria Operations. On top of the Home page is the search bar. The search bar is available on all pages of VMware Cloud Foundation Operations. To know more about the search capabilities, see the topic, Enhanced Search Capability.

The Launchpad helps you quickly get started with configuring and using VMware Cloud Foundation Operations. The Launchpad replaces the existing Quick Start page and provides unified use-case based, easy-to-follow workflows in the UI, to represent key supported capabilities of VMware Cloud Foundation Operations.

The Launchpad page is divided into the following sections:

**Pillars of Operations**

Helps you start a workflow based on the key operative areas of VMware Cloud Foundation Operations. Displays the following cards:

- Observability
- Capacity
- Cost
- Compliance
- Sustainability

When you click on any one of the cards, it takes you to a page that displays ways in which you can leverage that pillar of operation in VMware Cloud Foundation Operations. When you are on any of these pages, click **Learn More** in the top banner to understand the features of the product available to you for that pillar of operation and to understand any prerequisite steps.

**Applications**

Helps you manage your business applications and applications. Displays the following cards:

- Business Applications
- Applications

**Infrastructure**

Helps you monitor and evaluate the data based on your environment type, identifying trends in object behavior, calculating possible problems, and future capacity for objects in your system based on these trends. It displays alerts when an object exhibits defined symptoms. Displays the following cards:

- vSphere
- vSAN
- NSX
- Horizon
- Kubernetes

When you click on any one of the cards, it takes you to a page that displays ways in which you configure, manage and use the adapters. When you are on any of these pages, click **Learn More** to understand how to configure and use the adapter.

**VMware Cloud**

Helps you monitor and evaluate the data based on your VMware Cloud environment type. It identifies trends in object behavior, calculating possible problems and future capacity for objects in your system, based on the trends. Alerts you when an object exhibits the defined symptoms. Displays the following cards:

- VMware Cloud Foundation
- VMware Cloud on AWS
- Azure VMware Solution
- Google Cloud VMware Engine
- Oracle Cloud VMware Solution
- VMware Cloud on Dell EMC

When you click on any one of the cards, it takes you to a page that displays ways in which you configure, manage and use the adapters. When you are on any of these pages, click **Learn More** to understand how to configure and use the cloud adapter.

**Public Cloud**

Helps you monitor and evaluate the data based on your VMware Cloud environment type, identifying trends in object behavior. It calculates possible problems and future capacity for objects in your system based on the trends. Alerts you when an object exhibits defined symptoms. Displays the following cards:

- Amazon Web Services

- Microsoft Azure
- Google Cloud Provider

Click on any of the public cloud adapter cards to configure the cloud account if you already have not done so.

**Integrations**

Displays the following cards:

- Automation Central
- Outbound Plugins
- Content Repository
- Developer Center

## Installing Cloud Proxy

Using cloud proxies in VMware Aria OperationsVMware Cloud Foundation Operations, you can collect and monitor data from your remote data centers.

You can deploy one or more cloud proxies in VMware Cloud Foundation Operations to create a one-way communication between your end-point environment and VMware Cloud Foundation Operations

You can deploy one or more cloud proxies in VMware Aria Operations to create a one-way communication between your remote environment and VMware Aria Operations.

Log in to VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations, click **Data Sources › Cloud Proxies**.

| Task | Topics |
|---|---|
| Configure Cloud Proxies in VMware Cloud Foundation Operations | Configuring Cloud Proxies in VMware Aria Operations<br><br>See 'Configuring Cloud Proxies in VMware Aria Operations' in the *VMware Aria Operations Configuration Guide*. |
| Add a new Cloud Proxy in VMware Cloud Foundation OperationsVMware Aria Operations | Adding Cloud Proxy in VMware Aria Operations<br><br>See 'Adding Cloud Proxies in VMware Aria Operations' in the *VMware Aria Operations Configuration Guide*. |
| Activate Data Persistence in Clour Proxies | Data Persistence in Cloud Proxy.<br><br>See 'Data Persistence in Cloud Proxy' in the *VMware Aria Operations Configuration Guide* |
| Configure Cloud Proxies in AWS and Azure | Configuring Cloud Proxies in AWS and Azure<br><br>See 'Configuring Cloud Proxies in AWS and Azure' in the *VMware Aria Operations Configuration Guide*. |
| Monitor the Health of Cloud Proxies | Monitoring the Health of Cloud Proxies<br><br>See 'Monitoring the Health of Cloud Proxies' in the *VMware Aria Operations Configuration Guide*. |

*Table continued on next page*

*Continued from previous page*

| Task | Topics |
|---|---|
| Delete Cloud Proxies | Removing Cloud Proxies.<br><br>See 'Removing Cloud Proxies' in the *VMware Aria Operations Configuration Guide* |
| Upgrade Cloud Proxies | Upgrading Cloud Proxy.<br><br>See 'Upgrading Cloud Proxy' in the *VMware Aria Operations Configuration Guide* |
| Use the Cloud Proxy Command-Line Interface to<br>• Manually upgrade cloud proxy<br>• Generate support bundle<br>• Generate cloud proxy health status and connectivity details | Using the Cloud Proxy Command-Line Interface.<br><br>See 'Using the Cloud Proxy Command-Line Interface' in the *VMware Aria Operations Configuration Guide* |
| Using APIs with VMware Aria Operations Cloud Proxy | Using APIs with VMware Aria Operations Cloud Proxy.<br><br>See 'Using APIs with VMware Aria Operations Cloud Proxy' in the *VMware Aria Operations Configuration Guide* |
| Configure Collector Groups | Configuring Collector Groups.<br><br>See 'Configuring Collector Groups' in the *VMware Aria Operations Configuration Guide* |

## Monitoring Multiple Cloud Accounts in VMware Aria Operations

You can now monitor all your VMware cloud operations from the Overview page in VMware Aria Operations. This overview page provides a comprehensive understanding of your VMware cloud accounts which includes their geo-location, inventory, appliance health and management, and workload operations. You can also view the alerts and the cost for your entire VMware Cloud Infrastructure or individual VMware clouds.

The overview page provides a consolidated view of all VMware cloud accounts on the **All VMware Clouds** page. The different tiles provide insights into the location, health, cost, and inventory of your cloud accounts. When you first login, only vCenter, VMware Cloud Foundation, and Google Cloud VMware Engine cards are visible.

> **NOTE**
> No data displays for cloud accounts that do not have a configured account. To view data for the other VMware clouds or public clouds, you must configure cloud accounts from the **Administration › Integrations** page. For more information see the 'Adding Accounts' topic in the *VMware Aria Operations Configuration Guide*.

When you add any cloud account such as VMware Cloud on AWS, Azure VMware Solution, or Oracle Cloud VMware Solution hyperscale accounts, their respective cards will appear on the Overview page. This also applies to public clouds like AWS, Microsoft Azure, and Google Cloud Platform.

**World Map** tile - Use the world map to view the total number of locations and the various locations of your configured cloud accounts. The map shows numbers at various locations that signify groups of accounts. Multiple accounts in each area are grouped to provide a consolidated look. You can hover over the numbers to view the area that the grouped accounts belong to or click the number to zoom in and view accounts in that specific area. You can click an account to open the account summary tab. For more information see the 'Summary Tab' topic in the *VMware Aria Operations User Guide*.

**NOTE**
If multiple accounts of the same account type use the same region, it is considered as a single location.

**Inventory** tile - View the total number of configured accounts under each cloud account type and view the summarized resources count (vCenter, Domain, Datacenter, Cluster, Host, VM, Datastore, vSphere Distributed Switch, vSphere Distributed Port Group) for the whole cloud infrastructure.

**Appliances Health & Management** tile - Use the diagnostic findings, configuration drifts, and certificates of your VMware Clouds to track the health of your cloud accounts and manage their health.

**Workload Operations**: View the overall capacity, capacity remaining, and the time remaining before the capacity runs out. If no data displays as part of the capacity and time remaining sections, check your account configuration, and make sure they are collecting data.

**Cost**: View the total cost of ownership, potential savings, and realized savings. The amount is calculated from the first of the month to the current date. This value resets at the beginning of every month.

**Alerts** tile - View the total number of alerts need attention. Alerts of all cloud accounts are consolidated and represented graphically dating back a month from the current date dynamically.

**Sustainability** tile - Displays the sustainability status of your VMware Cloud that includes its green score, power consumption, carbon footprint, and overall environmental impact. To view the sustainability information in more detail, click **View Sustainability**. For more information, see the 'Sustainability' topic in the *VMware Aria Operations Configuration Guide*.

**Compliance** tile: View the compliance scores and misconfigurations of your cloud accounts.

You can select the individual VMware or public cloud card to view the account details. The tiles update to reflect values associated with that account.

**VMware Cloud - Example: vCenter Cloud Account**

The values displayed in the individual VMware cloud accounts like vCenter, VMware Cloud on Dell EMC, VMware Cloud on AWS, VMware Cloud Foundation, Azure VMware Solution, Oracle Cloud VMware Solution, and Google Cloud VMware Engine are based on their total number of accounts.

**NOTE**
To view the geo location of the accounts configured for vCenter and VMware Cloud Foundation, you must assign them to a physical data center. For more information, see the 'Adding Physical Data Centers' topic in the *VMware Aria Operations Configuration Guide*.

**Inventory** tile - Displays the top selected object types (vCenter, Datacenter, Cluster, Host, VM, Data Store, vSphere Distributed Switch, vSphere Distributed Port Group) hierarchically. Unspecified object types are not reflected in the tile.

**Top Objects Growth Trend** tile - Displays the growth trend of the top objects dating back a month from the current date dynamically.

**Public Cloud - Example: Amazon Web Services (AWS)**

The values displayed in the individual public cloud accounts like AWS, Microsoft Azure, Google Cloud Platform are based on their total number of accounts.

**Top Services** - Displays the list of top ten services for the selected cloud account. The list appears in descending order based on the count of each service.

**Top Services Growth Trend** - Displays the growth trend of the top services dating back a month from the current date.

## Secure the VMware Aria Operations Console

Secure the Console
After you install VMware Aria Operations, you secure the console of each node in the cluster by logging in for the first time.

1. Locate the node console in vCenter or by direct access. In vCenter, use Alt+F1 to access the login prompt.

   For security, VMware Aria Operations remote terminal sessions are deactivated by default.

2. Log in as `root`.
   VMware Aria Operations prevents you from accessing the command prompt until you create a root password.
3. When prompted for a password, press Enter.
4. When prompted for the old password, press Enter.
5. When prompted for the new password, enter the root password that you want, and note it for future reference.
6. Re-enter the root password.
7. Log out of the console.

## Log in to a Remote VMware Aria Operations Console Session

Log in to a Remote Console Session
As part of managing or maintaining the nodes in your VMware Aria Operations cluster, you might need to log in to a VMware Aria Operations node through a remote console.

For security, remote login is deactivated in VMware Aria Operations by default. To activate remote login, perform the following steps.

1. Log in to a vCenter Server system using a vSphere Web Client and select a vCenter Server instance in the vSphere Web Client navigator.
   a) Find the **Virtual Machine** in the hierarchy and click **Launch Console**.

      **NOTE**
      You can also use the vSphere Client to launch the node console by direct access after activating the SSHD service.

   The virtual machine console opens in a new tab of the Web browser.
2. Locate the node console and click **Launch Console**.
3. In vCenter, use Alt+F1 to access the login prompt and log in as `root`. If this is the first time logging in, you must set a root password.
   a) When prompted for a password, press Enter.
   b) When prompted for the old password, press Enter.
   c) When prompted for the new password, enter the root password that you want, and note it for future reference.
   d) Re-enter the root password.
4. To activate remote login, enter the following command:

   ```
   service sshd start
   ```

## Upgrade, Backup and Restore

You can update your existing VMware Aria Operations deployments to a newly released version.

When you perform a software update, you need to make sure you use the correct PAK file for your cluster. A good practice is to take a snapshot of the cluster before you update the software, but you must remember to delete the snapshot once the update is complete.

> **NOTE**
> During the upgrade process, any user modifications made to the default content, including alert definitions, symptoms, recommendations, policy definitions, views, dashboards, widgets, and reports, will be overwritten. To preserve such customizations and facilitate easy restoration after the upgrade, you must clone, export or create a backup of your content.

Starting with version 8.6 of VMware Aria Operations, internal certificates are renewed when you upgrade a cluster, except when the cloud proxy version 8.4, 8.5, or earlier is present. Automatic root-CA certificate renewal will be available when cloud proxy is version 8.6 and is upgraded to higher versions. After each product upgrade, the cluster will have a new root-CA certificate with a 5-year validity period.

> **NOTE**
> Automatic certificate renewal does not affect custom certificates.

## Obtain the Software Update PAK File

Each type of cluster update requires a specific PAK file. Make sure you are using the correct one.

**Download the Correct PAK files**

To update your VMware Aria Operations environment, you need to download the right PAK file for the clusters you wish to upgrade. In case modifications are required, you can manually update the hosts file after completing the software update.

To download the PAK file for VMware Aria Operations, go to Download VMware Aria Operations page and select the correct version from the drop-down list.

If you are using cloud proxy, download the *VMware Aria Operations Manager - Virtual Appliance upgrade .pak file with Cloud Proxy* file from the Product Downloads tab, to update the VMware Aria Operations environment and your cloud proxy together.

## Create a Snapshot as Part of an Update

It is mandatory to create a snapshot of each node in a cluster before you update a VMware Aria Operations cluster. Once the update is complete, you must delete the snapshot to avoid performance degradation.

For more information about snapshots, see the vSphere Virtual Machine Administration documentation.

1. Log into the VMware Aria Operations Administrator interface at `https://<primary-node-FQDN-or-IP-address>/admin`.
2. Click **Take Offline** under the cluster status.
3. When all nodes are offline, open the vSphere client.
4. Right-click a VMware Aria Operations virtual machine.
5. Click **Snapshot** and then click **Take Snapshot**.
   a) Name the snapshot. Use a meaningful name such as "Pre-Update."
   b) Uncheck the **Snapshot the Virtual Machine Memory** check box.
   c) Uncheck the **Ensure Quiesce Guest File System (Needs VMware Tools installed)** check box.
   d) Click **OK**.
6. Repeat these steps for each node in the cluster.

Start the update process as described in Install a Software Update.

## How To Preserve Customized Content

How To Preserve Customized Content
How To Preserve Customized Content in VMware Aria Operations
When you upgrade VMware Aria Operations, it is important that you upgrade the current versions of content types that allow you to alert on and monitor the objects in your environment. With upgraded alert definitions, symptom definitions, and recommendations, you can alert on the various states of objects in your environment and identify a wider range of problem types. With upgraded views, you can create dashboards and reports to easily identify and report on problems in your environment.

You previously customized versions of your alert definitions, symptom definitions, recommendations, or views.

You might need to perform certain steps before you upgrade the alert definitions, symptom definitions, recommendations, and views in your VMware Aria Operations environment.

- If you customized any of the alert definitions, symptom definitions, recommendations, or views that were provided with previous versions of VMware Aria Operations, and you want to retain those customized versions, perform the steps in this procedure.
- If you did not customize any of the alert definitions, symptom definitions, recommendations, or views that were provided with previous versions of VMware Aria Operations, you do not need to back them up first. Instead, you can start the upgrade, and during the upgrade select the check box named **Reset out-of-the-box content**.

1. Before you begin the upgrade to VMware Aria Operations, back up the changes to your alert definitions, symptom definitions, recommendations, and views by cloning them.
2. Start the upgrade of VMware Aria Operations.
3. During the upgrade, select the check box named **Reset out-of-the-box content**.

After the upgrade completes, you have preserved your customized versions of alert definitions, symptom definitions, recommendations, and views, and you have the current versions that were installed during the upgrade.

Review the changes in the upgraded alert definitions, symptom definitions, recommendations, and views. Then, determine whether to keep your previously modified versions, or to use the upgraded versions. For more information, see Creating a Backup and Importing Content in the Managing Content chapter of the Configuration Guide.

## Back Up and Restore

Back up and restore your VMware Aria Operations system regularly to avoid downtime and data loss in case of a system failure. If your system does fail, you can restore the system to the last full or incremental backup.

You can back up and restore VMware Aria Operations single or multi-node clusters by using vSphere Data Protection or other backup tools. You can perform full, differential, and incremental backups and restores of virtual machines.

To back up and restore VMware Aria Operations single or multi-node clusters using the Veeam Backup & Replication tool, see About Veeam Backup & Replication.

It is highly recommended to take a backup during quiet periods. Since a snapshot based backup happens at the block level, it is important that there are limited or no changes being performed by a user on the cluster configuration. This will ensure that you have a healthy backup.

It is best to take the cluster offline before you back up the VMware Aria Operations nodes. This will ensure the data consistency across the nodes and internally in the node. You can either shut down the VM before the backup or activate quiescing.

If the cluster remains online, backup your VMware Aria Operations multi-node cluster by using vSphere Data Protection or other backup tools, deactivate quiescing of the file system.

**NOTE**
All nodes are backed up and restored at the same time. You cannot back up and restore individual nodes.

## VMware Aria Operations Software Updates

Software Updates
VMware Aria Operations includes a central page where you can manage updates to the product software.

### How Software Updates Work

The Software Update option lets you install updates to the VMware Aria Operations product itself.

### Where You Find Software Updates

Log in to the VMware Aria Operations administration interface at https://*primary-node-name-or-ip-address*/admin. On the left, click **Software Update**.

### Software Update Options

The options include a wizard for locating the update PAK file and starting the installation, plus a list of updates and the VMware Aria Operations cluster nodes on which they are installed.

**Table 22: Software Update Options**

| Option | Description |
|---|---|
| Install a Software Update | Launch a wizard that allows you to locate, accept the license, and start the installation of a VMware Aria Operations software update. |
| Node Name | Machine name of the node where the update is installed |
| Node IP Address | Internet protocol (IP) address of the node where the update is installed. Primary and replica nodes require static IP addresses. Data nodes may use DHCP or static IP. |
| Update Step | Software update progress in step x of y format |
| Status | Success, failure, in-progress, or unknown condition of the software update.<br><br>For cloud proxy upgrade, every stage of the upgrade process is displayed. Hover the mouse near the status message to see more details in the pop-up window. The Cloud Proxy upgrade stages are as follows:<br>• Stage 1 - Downloading<br>• Stage 2: Extracting<br>• Stage 3: Upgrading<br>• Stage 4: Rebooting<br>• Stage 5: Success |

### Install a Software Update

If you have already installed VMware Aria Operations, you can update your software when a newer version becomes available.

- Create a snapshot of each node in your cluster. For information about how to perform this task, see the VMware Aria Operations Information Center.
- Create a snapshot of each node in your cluster. See Create a Snapshot as Part of an Update for details.

- Obtain the PAK file for your cluster. For information about which file to use, see the VMware Aria Operations Information Center.
- Obtain the PAK file for your cluster. See Obtain the Software Update PAK File for details.
- Before you install the PAK file, or upgrade your VMware Aria Operations instance, clone any cusomized content to preserve it. Customized content can include alert definitios, symptom definitions, recommendations, and views. Then, during the softeare update, you select the option named **Install the PAK file even if it is already installed**.
- Since version 6.2.1, VMware Aria Operations update operation has a validation process that identifies issues before you start to update your software. Although it is good practice to run the pre-update check and resolve any issues found, users who have environmental constraints can deactivate this validation check.
  To deactivate the pre-update validation check, perform the following steps:
  - Edit the update file to `/storage/db/pakRepoLocal/bypass_prechecks_VMware AriaOperationsManagerEnterprise-buildnumberofupdate.json`.
  - Change the value to TRUE and run the update.

    **NOTE**
    If you deactivate the validation, you might encounter blocking failures during the update itself.

    **NOTE**
    Installation might take several minutes or even a couple hours depending on the size and type of your clusters and nodes.

    **NOTE**
    Before you can upgrade to VMware Aria Operations 8.16 you must delete all remote collectors and replace them with cloud proxies. For information about migrating from VMware Aria Application Remote Collector to cloud proxy, see KB 83059.

1. Log into the primary node VMware Aria Operations administrator interface of your cluster at `https://primary-node-FQDN-or-IP-address/admin`.
2. Click **Software Update** in the left pane.
3. Click **Install a Software Update** in the main pane.
4. Follow the steps in the wizard to locate and install your PAK file.

   This updates the OS on the virtual appliance and restarts each virtual machine.
5. Read the **End User License Agreement** and **Update Information**, and click **Next**.
6. Click **Install** to complete the installation of software update.

   **NOTE**
   After you click **Install**, the installer will restart the VMware Aria Operations administrator interface, and you will be logged out. Log in once again to the VMware Aria Operations administrator interface when it is ready, and follow the update status in the software update page.
7. Log back into the primary node administrator interface.
   The main Cluster Status page appears and cluster goes online automatically. The status page also displays the Bring Online button, but do not click it.
8. Clear the browser caches and if the browser page does not refresh automatically, refresh the page.
   The cluster status changes to Going Online. When the cluster status changes to Online, the upgrade is complete.

   **NOTE**
   If a cluster fails and the status changes to offline during the installation process of a PAK file update, then some nodes become unavailable. To fix this, you can access the administrator interface and manually take the cluster offline and click **Finish Installation** to continue the installation process.

9. Click **Software Update** to check that the update is done.
   A message indicating that the update completed successfully appears in the main pane.

   **NOTE**
   When you update VMware Aria Operations to a latest version, all nodes get upgraded by default.

   If you are using cloud proxies, the cloud proxy upgrades start after the VMware Aria Operations upgrade is complete successfully. For more information, see the Monitoring the Health of Cloud Proxies from the Admin UI topic in the *VMware Aria Operations Configuration Guide*.

Delete the snapshots you made before the software update.

**NOTE**
Multiple snapshots can degrade performance, so delete your pre-update snapshots after the software update completes.

## Before Upgrading to VMware Aria Operations 8.18

With every VMware Aria Operations release, many metrics are either discontinued or deactivated. These changes update the capacity analytics and improve the product scale. VMware has made many of these changes transparent or nearly so. Still, multiple changes can impact management packs that you might be using, along with the dashboards and reports that you have created. Therefore, before upgrading, run the VMware Aria Operations Pre-upgrade Readiness Assessment Tool (Assessment Tool) that helps you understand the precise impact on your environment through a detailed report.

### Why Run the Assessment Tool

Various changes in VMware Aria Operations can impact the user experience. When you run the Assessment Tool, you get an HTML-formatted report identifying all the points in your system affected by the changes. Further, the Assessment Tool gives recommendations for the correct changes to be made in your content for when you upgrade from a previous release.

**NOTE**
You must run the Assessment Tool on the instance of the VMware Aria Operations installation that you want to assess - typically your production system. The Assessment Tool does not alter anything in your system, and deletes itself when it has completed its run. It leaves behind only the assessment result - a support bundle that you download from the Support Bundles section of the VMware Aria Operations Administration user interface. The Assessment tool validates your environment to ensure it is ready for the upgrade. For example, if the ESXi version does not match the product requirements, the assessment tool will identify the issue and provide you with a recommendation in the Systems Validation tab.

For detailed instructions on running the Assessment Tool, see Running the 8.18 Pre-Upgrade Readiness Assessment Tool.

To view the upgrade path from an earlier version of VMware Aria Operations to 8.18, see VMware Aria Operations Upgrade Path.

#### Running the VMware Aria Operations 8.18 Pre-Upgrade Readiness Assessment Tool

Before upgrading, you can gauge the impact on your system by running the VMware Aria Operations Pre-Upgrade Readiness Assessment Tool (Assessment Tool). The tool generates a report detailing the precise impact on your environment and gives suggestions for replacement metrics.

You must have administrator privileges in your current installation of VMware Aria Operations to download and run the Assessment Tool. For more information on using the upgrade assessment tool, see the following KB article 67311.

Using the Assessment Tool consists of four distinct steps:

1. Download the corresponding version of the PAK file from Broadcom Support Portal.
2. Run the VMware Aria Operations Pre-Upgrade Readiness Assessment Tool.
3. Extract the report from the generated ZIP file.
4. Click the various items in the report to link to the solutions grid.

**NOTE**
You must run the Assessment Tool on the instance of the VMware Aria Operations installation that you want to assess - typically your production system. The Assessment Tool does not alter anything in your system, and deletes itself when it has completed its run. It leaves behind only the assessment result - a support bundle that you download from the Support Bundles section of the VMware Aria Operations Administration user interface.

1. Download the corresponding version of the Assessment Tool PAK from Broadcom Support Portal to your local machine. Search for APUAT or VMware Aria Operations - Upgrade Assessment Tool.
2. Open a browser and navigate to the VMware Aria Operations administrator console: https://*<primary_node_IP>*/admin.
   Then log into the administrator user interface with the user ID **admin** and the associated password.
3. In the left pane of the administration home page, click **Software Update**.
   The Software Update screen appears.
4. Click **Install a Software Update** at the top of the screen.
   The Add Software Update workspace appears.
5. Click the **Browse** link and navigate to the PAK file you downloaded in Step 1.
   A check mark appears next to the statement: **The selected file is ready to upload and install. Click UPLOAD to continue.**
6. Ensure that a check mark appears next to the statement: **Install the PAK file even if it is already installed**.
   Leave blank the check box next to Reset Default Content...
7. Click the **UPLOAD** link.
   The PAK file is uploaded from your local machine to VMware Aria Operations. Uploading may take a few minutes.
8. Once the PAK file is uploaded, click **NEXT**.
   The End User License Agreement appears.
9. Click the check box next to the statement: **I accept the terms of this agreement**.
   Click **NEXT**. The Important Update and Release Information screen appears.
10. Review the release information and click **NEXT**. At the Install Software Update screen, click **INSTALL**.
    The Software Update screen appears again, this time with a rotating icon and an **installation in progress...** bar marking the progress of the PAK file and assessment as they run on your environment. The process can take from five to 20 minutes, depending on the size of your system.
11. When the process is complete, click **Support** in the left pane.
    The Support screen appears.
12. Select the **Support Bundles** option above the toolbar.
    The available support bundles are listed.
13. Locate the support bundle most recently created. Click the chevron next to the bundle name to open the file and select it, then click the download link on the toolbar to save the support bundle ZIP file to your local files.
14. To review the report, extract the files from the ZIP file and open the HTML file. (Do not open the CSV file, it is for VMware use only.)
    The report is a graphical depiction of your VMware Aria Operations UI components - dashboards, reports, management packs, alerts, heat maps, and so on - and includes the number of deprecated metrics impacting each component. For example, you might find that 10 of your 25 dashboards contain a total of 15 deprecated metrics.
15. Click a component.

The report details for that component are listed following the graphics, under Impacted Component Details. Taking dashboards as an example, the list provides - for each dashboard - the dashboard name, owner, widgets removed, metric-impacted views, and metric-impacted widgets. The deprecated metrics are live links.

16. Click a live metric link.
    A browser window opens at URL http://partnerweb.vmware.com/programs/vrops/DeprecatedContent.html with the selected metric highlighted in a table of like metrics. If a replacement metric is available for the deprecated metric, it is listed in the same row by name and metric key. You might choose to install the new metric in place of the deprecated metric.
17. Repeat Steps 15 and 16 for all your components.
    If you replace the deprecated metrics with new metrics, or update each component to provide needed information without the deprecated metrics, your system is ready for the upgrade.
18. Rerun the entire assessment process from Step 1 to confirm that your system is no longer impacted or at least mostly not impacted by the metrics changes.
19. Once you have upgraded to VMware Aria Operations 8.18, fix the remaining issues with replacement metrics available in the new release.

Your VMware Aria Operations components are updated to work correctly in the 8.18 release.

Once you have installed VMware Aria Operations 8.18, conduct, at a minimum, random testing to determine if system metrics are operating as you expect. Monitor the platform on an ongoing basis to confirm that you are receiving the correct data.

# VMware Aria Operations Configuration Guide (8.18)

The *Configuring VMware Aria OperationsVMware Cloud Foundation Operations* guide describes how to configure and monitor your environment. It shows you how to connect VMware Aria OperationsVMware Cloud Foundation Operations to external data sources and analyze the data collected from them, ensure that users and their supporting infrastructure are in place, configure resources to determine the behavior of your objects, and format the content that appears in VMware Aria OperationsVMware Cloud Foundation Operations.

## Intended Audience

This information is intended fo administrators, virtual infrastructure administrators, and operations engineers who configure, monitor, manage, and maintain the objects in your environment.

For users who want to configure VMware Aria OperationsVMware Cloud Foundation Operations programmatically, the VMware Aria OperationsVMware Cloud Foundation Operations REST API documentation is available in HTML format and is installed with your VMware Aria OperationsVMware Cloud Foundation Operations instance. For example, if the URL of your instance is `https://VMwareAriaOps.example.com`, the API reference is available from `https://VMwareAriaOps.example.com/suite-api/docs/rest/index.html`.

## Configuring VMware Aria Operations

The *Configuring VMware Aria OperationsVMware Cloud Foundation Operations* guide describes how to configure and monitor your environment. It shows you how to connect VMware Aria OperationsVMware Cloud Foundation Operations to external data sources and analyze the data collected from them, ensure that users and their supporting infrastructure are in place, configure resources to determine the behavior of your objects, and format the content that appears in VMware Aria OperationsVMware Cloud Foundation Operations.

To help you maintain and expand your VMware Aria Operations installation, this information describes how to manage nodes and clusters, configure NTP, view log files, create support bundles, and add a maintenance schedule. It provides information about license keys and groups, and shows you how to generate a passphrase, review the certificates used for authentication, run the describe process, and perform advanced maintenance functions.

## Intended Audience

This information is intended for VMware Aria OperationsVMware Cloud Foundation Operations administrators, virtual infrastructure administrators, and operations engineers who install, configure, monitor, manage, and maintain the objects in your environment.

For users who want to configure VMware Aria OperationsVMware Cloud Foundation Operations programmatically, the VMware Aria OperationsVMware Cloud Foundation Operations REST API documentation is available in HTML format and is installed with your VMware Aria OperationsVMware Cloud Foundation Operations instance. For example, if the URL of your instance is `https://VMwareAriaOps.example.com`, the API reference is available from `https://VMwareAriaOps.example.com/suite-api/docs/rest/index.html`.

## About Configuring VMware Aria OperationsVMware Cloud Foundation Operations

You configure objects, alerts, actions, policies, dashboards, and reports, in VMware Aria OperationsVMware Cloud Foundation Operations to effectively monitor your environment. You use administration settings to manage your environment.

Configure solutions in VMware Aria OperationsVMware Cloud Foundation Operations to connect to and analyze data from external data sources in your environment. Once connected, you use VMware Aria OperationsVMware Cloud Foundation Operations to monitor and manage objects in your environment. Solutions that are installed together with VMware Aria OperationsVMware Cloud Foundation Operations include vSphere, Logs, VMware Aria Automation, VMware vSAN, and many more. Configure these adapters to connect to and integrate with these instances.

Create alert definitions so that whenever there is a problem, VMware Aria OperationsVMware Cloud Foundation Operations triggers alerts and provides recommendations to resolve the problem. The process of configuring alerts involves defining alerts, symptoms, and recommendations.

Activate actions to address a problem in the monitored environment. The actions let you resolve a problem by remaining in the VMware Aria OperationsVMware Cloud Foundation Operations environment itself.

Create a policy to define rules for VMware Aria OperationsVMware Cloud Foundation Operations to use. You can use a policy to analyze and display information about the objects in your environment.

Define compliance standards to determine the compliance of your objects. You can use VMware Aria OperationsVMware Cloud Foundation Operations alert definitions to create compliance standards that notify you when an object does not comply with a required standard.

Create super metrics to give you a big picture of your environment. A super metric is a mathematical formula that contains one or more metrics. It is a custom metric that you design and is useful when you need to track combinations of metrics, either from a single object or from multiple objects. If a single metric cannot tell you what you need to know about the behavior of your environment, you can define a super metric.

Create dashboards to determine the nature and timeframe of existing and potential issues with your environment. You create dashboards by adding widgets to a dashboard and configuring them.

Create views to interpret metrics, properties, and policies of various monitored objects including alerts. Generate a report to capture details related to current or predicted resource needs. A report is a scheduled snapshot of views and dashboards.

## Accessibility Compliance

VMware Cloud Foundation OperationsVMware Aria Operations accessibility compliance provides several interactive elements that can be operated using a keyboard and screen reader.

## Keyboard Support

**Table 23: Tooltips, Grid Sorting, Drag and Drop, and Combo-box with the X icon**

| Component | Description | Examples |
|---|---|---|
| Open and Close Tooltips | Navigate through elements using the **TAB** key.<br>Open tooltips using the **Ctrl + i** keys.<br><br>Close tooltips using the **ESC** key. | Navigate through elements in the workbench page and open and close tooltips.<br>1. From the left menu, click **Operations** > **Troubleshoot**.<br>2. Click on a card.<br>If no card is available, search for a resource and click on it.<br>3. Use the **TAB** key to navigate through the elements.<br>4. Click **Ctrl + i** to open a tooltip and once done, click the **ESC** key to close it. |
| | | Navigate to the object relation chart and open and close tooltips.<br>1. From the left menu, navigate to **Inventory** > **Inventory Panel** |

*Table continued on next page*

*Continued from previous page*

| Component | Description | Examples |
|---|---|---|
| | | **(Detailed View)** > **Integrations** > **All Objects**.<br>2. Use the hierarchies in the left pane to locate the objects that you want and then click the **Metrics** tab.<br>3. Click **Show Object Relationship**.<br>4. Use the **TAB** key to navigate through the elements.<br>5. Click **Ctrl + i** to open a tooltip and once done, click the **ESC** key to close it. |
| Open Tooltips | Navigate through the alerts grid using the **TAB** key.<br>Open tooltips using the **Ctrl + i** keys. | Navigate through the inventory page and open tooltips.<br>1. From the left menu, click **Operations** > **Configurations**, and then click the **Inventory Management** tile > **Objects** tab.<br>2. Use the **TAB** key to navigate to the **Relevance** column and then click **Ctrl+i** to open the tooltip. |
| | | Navigate through the alerts grid and open tooltips.<br>1. From the left the menu, click **Operations** > **Alerts**.<br>2. Select an alert from the list to enable the Actions menu.<br>3. Use the **TAB** key to navigate to the **Importance** column and then click **Ctrl+i** to open the tooltip. |
| Grid Sorting | Sort columns that can be sorted using the **Enter** or **Space** keys. | Sort a grid.<br>1. Navigate to a column header.<br>2. Use the **Enter** or **Space** keys to sort the columns. |
| Drag and Drop | Drag and drop elements using the **TAB** and **Enter** keys.<br>**NOTE**<br>If the default functionality of the **Enter** key has changed, you must use **Ctrl+Enter** instead. | Drag and drop alert symptoms.<br>1. From the left the menu, click **Operations** > **Configurations**, and then click the **Alert Definitions** tile.<br>2. Click **Add** and enter the alert definition details and then click **Next**.<br>3. On the **Symptoms** tab, use the **TAB** key to navigate through the grid and click the **Enter** key on the first column to select one of the symptoms. |

*Table continued on next page*

*Continued from previous page*

| Component | Description | Examples |
|---|---|---|
| | | 4. Use the **TAB** key again to navigate through the drop areas and then click the **Enter** key to drop the symptom.<br>5. Click **ESC** to cancel the action. |
| Combo-box with an X icon | Use the **X** icon or the **Delete** key to clear any combo-box in VMware Cloud Foundation OperationsVMware Aria Operations. | Clear the combo-box for alerts.<br>1. From the left menu, click **Operations** > **Configurations**, and then click the **Alerts Definitions** tile.<br>2. Click **Add** and enter the alert definition details and then click **Next**.<br>3. Click the **X** icon to clear it.<br>4. Click the **Delete** key to clear it. |

## Collecting Data with Cloud Proxies in VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations

Collecting Data with Cloud Proxies
You can use cloud proxies in VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations to collect and monitor data from your on-premises data centers across different geo locations.

**NOTE**
FIPS mode is supported in cloud proxy. To leverage this functionality, make sure your cluster is in FIPS mode.

Cloud proxies provide high availability within your cloud environment, you can group two or more cloud proxies to form a collector group. The cloud proxy collector group ensures that there is no single point of failure in your cloud environment. If one of the cloud proxies experiences a network interruption or becomes unavailable, the other cloud proxy from the collector group takes charge and ensures that there is no downtime. All other user-initiated manual operations on the collector, such as to stop or restart the collector manually, do not result in automated rebalancing. For more information, see the 'Configuring Collector Groups' topic in the *Configuring VMware Aria Operations* guide.

**NOTE**
When cloud proxies provide high availability within your cloud environment, the cluster can survive the loss of one data node without losing any data. However, the cloud proxy does not guarantee that the adapter instance will collect all the data during adapter instance failover (or when reassigning the adapter instance). VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations cloud proxy only provides additional application level data protection to ensure application level availability.

You can also use cloud proxies to rebalance the resources across the collectors in your collector group. The Rebalance option is available as part of the Edit menu in the Collector Groups page.

**NOTE**
You can use the rebalance option before the vCenter Adapter initiates data collection. Once the data collection starts, the rebalance option is disabled.

## Configuring Cloud Proxies in VMware Aria Operations

Configuring Cloud Proxies

Using cloud proxies in VMware Aria Operations, you can collect and monitor data from your remote data centers. Typically, you need only one cloud proxy per physical data center. You can deploy one or more cloud proxies in VMware Aria Operations to create a one-way communication between your remote environment and VMware Aria Operations. Cloud proxies collect data from the end-point environment and uploads it to VMware Aria Operations. Cloud proxies can support multiple vCenter accounts. For more information on cloud proxies, see the topic called Cloud Proxy FAQ in the *Configuring VMware Aria Operations Guide*.

- Verify that you have an IP address, a DNS entry, and permissions to deploy OVF templates in vSphere.
- Log in to vSphere and verify that you are connected to a vCenter system.
- Allow outgoing HTTPS traffic for cloud proxy over port 443. For more information on firewall requirements in VMware Aria Operations see KB article 93210.
- Allow incoming traffic to cloud proxy over ports 443, 8443, 4505, and 4506 for telegraf based application monitoring.
- Allow incoming traffic to cloud proxy over port 443 for push model adapters or Suite-API on cloud proxy.
- Add a vCenter cloud account and provide an account with the following read and write privileges:
  - vCenter IP address or FQDN
  - Permissions required to install a cloud proxy on the vCenter Server.

  For more information on privileges, see the topic called "Privileges Required for Configuring a vCenter Adapter Instance" in the *VMware Aria Operations Configuration Guide*.
- Cloud proxies must have a proper DNS resolution to the VMware Aria Operations nodes when using short/long FQDN names. This is applicable to on-prem cloud proxy.
- Using a firewall to restrict traffic by IP is not recommended since IPs can change without notice. Restricting traffic must be performed via FQDNs only.

1. Log in to VMware Aria Operations.
2. From the left menu, click **Administration › Cloud Proxies**, and then click **Add**.
3. Click **Broadcom support portal** to navigate to the Broadcom support portal, log in using your credentials. and download the cloud proxy OVA file.
4. Navigate to your vSphere, select the name of your vCenter cluster, and select **Deploy OVF Template** from the **Actions** menu.
5. Insert the ova link and then click **Next**.
   - Paste the cloud proxy ova link in the **URL** field.
   - Click the **Local File** option, browse, and select the downloaded OVA file.
6. Follow the prompts to install the OVA on your vCenter.

   For the most current information about sizing and scaling, see Knowledge Base article 78491.
7. When prompted to enter the Unique Registration Key in the **Customize template** screen, return to the Install Cloud Proxy page in VMware Aria Operations.
8. Activate **Data Persistence** to store data in the cloud proxy in case of connectivity issues. For more information, see Activating Data Persistence in Cloud Proxy.
9. Activate **Log Forwarding** to use the cloud proxy for forwarding logs from end points to VMware Cloud Foundation Operations for logs.

   **NOTE**
   Log forwarding in cloud proxy is available only if you have configured a VMware Cloud Foundation Operations for logs account in VMware Aria Operations. For more information, see Configuring VMware Cloud Foundation Operations for logs with VMware Aria Operations.
   If you activate log forwarding in cloud proxy, the cloud proxy is used solely for log forwarding and cannot be used within collector groups.
10. Click the **Copy Key** icon.

The unique registration key expires 24 hours after generation. To avoid using an expired key, click **Regenerate Key** before proceeding. The unique registration key is used by the cloud proxy to authenticate to VMware Aria Operations.

**NOTE**
The unique registration key is refreshed and a new key is generated if you reload the cloud proxy page. A new unique registration key generated if you activate data persistence, or log forwarding, or both.

11. Return to vSphere and paste the key in the **Unique Registration Key** text box to install the VMware Aria Operations Cloud Appliance.
12. Select **Prefer IPv6** to use IPv6 for internal communications. For more information, see the topic, 'Using IPv6 with VMware Aria Operations' in the *Getting Started with VMware Aria Operations Guide*.
13. Set up a proxy server in the **Customize template** screen.
    a) Enter details in the **Network Proxy IP Address** and **Network Proxy Password** properties.

    **NOTE**
    If you use network proxy for log forwarding, port 9543 must be open.

    b) To activate SSL, select the **Use SSL connection to proxy** check box.
    c) If you are using SSL, you can verify the certificate of the proxy server. Public certificate authorities are used to verify the proxy server certificate. To activate this, select the **Verify proxy's SSL cert** check box in the **Verify SSL cert** property.
    d) You can specify the IP /FQDN URL that is used to access the system when a load balancer is used.
    e) If you have a custom certificate authority, paste the root certificate authority in the **Custom CA** property to verify the certificate of the proxy server. The root certificate authority is passed on to the cloud proxy. Include the following lines when you copy the root certificate authority:

    ```
    "-----BEGIN CERTIFICATE-----"


    "-----END CERTIFICATE-----"
    ```

14. Click **Finish**.

    The deployment takes a few minutes to finish.

15. Locate the cloud proxy you just installed, select the VMware Aria Operations Cloud Appliance, and click **Power on**.

    **NOTE**
    You must power on the VMware Aria Operations Cloud Appliance within 24 hours of registering it. After 24 hours, the Unique Registration Key expires, and you must delete the VMware Aria Operations Cloud Appliance and deploy another cloud proxy.

16. Return to the Cloud Proxy page in VMware Aria Operations to view the status of the cloud proxy you just installed. For more information, see Monitoring the Health of Cloud Proxies.

17. To view the accounts that are using this connection, click the Cloud Proxy.

    The communication from the cloud proxy to cloud is one way. The cloud proxy initiates this connection and if necessary, it also pulls data from cloud (like the adapters configuration or upgrade pak). The cloud proxy requires a regular Internet access over the HTTPS protocol but it does not need any special firewall configuration. The cloud proxy verifies the certificate of the cloud service it connects to and if there are transparent proxy servers which do stop SSL, it might cause connectivity problems for the cloud proxy.
    The cloud proxy also supports connection through the corporate proxy server. The proxy settings are given during OVF deployment.

18. To delete a cloud proxy, click the vertical ellipsis and then click **Delete**. For more information, see Deleting Cloud Proxies .

Upgrade your cloud proxy. For more information, see the topic called Upgrading Cloud Proxy in the *Getting Started with VMware Aria Operations Guide*.

The VMware vSphere solution connects VMware Aria Operations to one or more vCenter instances. For more information see the topic called Configure a vCenter Server Cloud Account in VMware Aria Operations in the *Configuring VMware Aria Operations Guide*.

### Activating Data Persistence in Cloud Proxy

Using data persistence, you can avoid data gaps in case of temporary connectivity issues. Activate data persistence to allow the cloud proxy to store data that is sent to VMware Aria OperationsVMware Cloud Foundation Operations, if the connection fails between the cloud proxy and VMware Aria OperationsVMware Cloud Foundation Operations. The cloud proxy stores all the persisted data that includes metrics and properties, along with time stamps.

Cloud proxy can store data for a maximum duration of one hour. If there is lack of space or if the connection fail lasts for more than an hour, the cloud proxy rotates the stored data by deleting the oldest stored data and replacing it with the most recently collected data. In case of lack of space, you can add additional storage, for more information, see KB article 2016022.

Once the connection is restored, the cloud proxy sends the stored data to VMware Aria OperationsVMware Cloud Foundation Operations before the real-time data. The stored data is displayed before the real-time data as data is displayed in the correct time series. If the connection loss was more than one hour, there can also be a gap in the displayed data.

You can activate data persistence when you deploy a new cloud proxy. For more information, see Configuring Cloud Proxies in . You can also activate data persistence in an existing cloud proxy. For more information, see Monitoring the Health of Cloud Proxies.

## Configuring Cloud Proxies in AWS and Azure

Using cloud proxies in VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations, you can collect and monitor data from your public cloud environments (Amazon Web Services and Microsoft Azure). You can deploy one or more cloud proxies on VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations to create a one-way communication between your end-point environment and VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations. The cloud proxies work as one-way remote collectors and upload data from the end-point environment to VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations. Cloud proxies can support multiple cloud accounts.

- Allow outgoing HTTPS traffic for cloud proxy using port 443.
- Allow outgoing traffic from the endpoints to cloud proxy using port 443.

1. Log in to VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations.
2. From the left menu, click **Administration › Cloud Proxy**, and then click **New**.
3. Click the **Target Location** drop-down and select one of the following.
   – **Amazon Web Services** (AWS).

      1. Click the **AWS image URL**.
         The AWS portal opens. Log in to your AWS account if prompted.
      2. Click **Launch Instance with AMI**. For more information on how to launch the instance from AMI, see the VMware KB article 88673.

3. Click the **Regenerate Key** icon to avoid using an expired key and then click the **Copy Key** icon to copy the unique registration key for the User Data field in the Advanced section on the AWS EC2 instance creation page.
4. Launch the instance. It will take a few minutes for VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations to detect the newly deployed cloud proxy after it is launched and powered up in AWS.

– **Microsoft Azure**.

1. Click the **Microsoft Azure URL**.
   The Microsoft Azure portal opens. Log in to your Microsoft Azure account if prompted.
2. Click **Create**. For more information on how to create a VM from Microsoft Azure, see the VMware KB article 88672.
3. Click the **Regenerate Key** icon to avoid using an expired key and then click the **Copy Key** icon to copy the unique registration key for the User Data field in the Advanced section on the Microsoft Azure VM instance creation page.
4. Launch the instance. It will take a few minutes for VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations to detect the newly deployed cloud proxy after it is launched and powered up in Microsoft Azure.

4. Return to the Cloud Proxy page in VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations to view the status of the cloud proxy you just installed. For more information, see Monitoring the Health of Cloud Proxies.

For AWS you must deploy one cloud proxy per account. To be able to collect metrics from different regions, you must ensure the connectivity between the AWS EC2 instances and cloud proxy.

For Microsoft Azure you must deploy one cloud proxy per account. To be able to collect metrics from different regions, you must ensure the connectivity between the Microsoft AzureVM instances and cloud proxy.

## Monitoring the Health of Cloud Proxies

You can view the status and health of your cloud proxy after you add it in VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations. You can then monitor the health and view alerts and metrics of your cloud proxy using the VMware Aria Operations Cloud Proxy object.

1. Log in to VMware Aria OperationsVMware Aria OperationsVMware Cloud Foundation Operations.
2. From the left menu, click **Data Sources › Cloud Proxies**.
   The list of cloud proxies is displayed.

| Option | Description |
|---|---|
| Name | The name of the cloud proxy. |
| IP | The IP address of the cloud proxy. |
| Status | Status of the cloud proxy. For example, the Going Online status is displayed for a few minutes when you add a new cloud proxy. Once the cloud proxy is connected to VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations, the status changes to Online. If the VMware Aria OperationsVMware Cloud Foundation |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | OperationsVMware Aria Operations is not connected, the Offline status is displayed. |
| Version | The version used to install the cloud proxy. |
| Accounts | The number of accounts that are created and associated with the cloud proxy. |
| Type | Displays the cloud proxy type. |
| Collector Group | Displays if the cloud proxy is part of the collector group. If the cloud proxy is part of the collector group, the name is displayed. Click the name to view the collector group details.<br><br>**NOTE**<br>No data is displayed if the cloud proxy is a standalone cloud proxy and not a part of the collector group. |
| Network Proxy Address | The network proxy address of the cloud proxy. |
| Network Proxy Port | The network proxy port number of the cloud proxy. |
| Target | Displays the target location on which the cloud proxy is deployed. |
| Data Persistence | Displays the status of data persistence for the cloud proxy. Data persistence activates the cloud proxy to store data if the connection fails between the cloud proxy and VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations.<br>• **Activated**. Cloud proxy will store data.<br>• **Deactivated**. Cloud proxy will not store data.<br>To activate/deactivate data persistence for multiple cloud proxies, select the cloud proxies, click the horizontal ellipsis, and then select **Activate Data Persistence** or **Deactivate Data Persistence**.<br><br>To activate/deactivate data persistence for a single cloud proxy, select the cloud proxy, click the vertical ellipsis, and then select **Activate Data Persistence** or **Deactivate Data Persistence**.<br><br>**NOTE**<br>When the connection is restored, the cloud proxy sends the stored data to VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations. The stored data is displayed before the real-time data is displayed. |
| Time Estimation | Displays the estimated time duration for which the cloud proxy persists data.<br>Cloud proxy can store data for a maximum duration of one hour. If there is a lack of space or if the connection fail lasts for more than an hour, the cloud proxy rotates |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | the stored data by deleting the oldest stored data and replacing it with the most recently collected data.<br><br>**NOTE**<br>The time estimate value is displayed only if data persistence is activated. |
| Filter | You can search the list of cloud proxies according to the following criteria:<br>• Name<br>• IP<br>• Status<br>• Version<br>• Collector Group<br>• Network Proxy Address<br>• Network Proxy Port<br>• Target<br>• Data Persistence |

3. Click the vertical ellipsis.
   a) Click **Activate/Deactivate Data Persistence** to activate or deactivate data persistence for the cloud proxy.
   b) Click **Remove** to delete the cloud proxy.
4. Click a **Cloud Proxy** name.
   The **Cloud Proxy Details** open.

**Table 24: Cloud Proxy Details**

| Option | Description |
|---|---|
| Rename | Click the **Rename Cloud Proxy** icon to update the cloud proxy name. |
| Proxy ID | ID of the cloud proxy. |
| IP Address | IP address of the cloud proxy. |
| OVA Version | The OVA file version used to install the cloud proxy. |
| Status | Status of the cloud proxy. For example, the Getting Online status is displayed for a few minutes when you add a cloud proxy. Once the cloud proxy is connected to VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations, the status changes to Online. If the VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations is not connected, the Offline status is displayed. |
| Creation Date | Date of creation of the cloud proxy. |
| Last Heartbeat | Last time stamp when VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations ran a Health Check for this cloud proxy. When you click a cloud proxy to view its details, VMware Aria OperationsVMware Cloud Foundation |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | OperationsVMware Aria Operations sends a heartbeat to check if the cloud proxy is still reachable. |
| CPU | CPU usage. |
| Memory | Memory usage. |
| Number of Cloud Accounts | Displays the number of cloud accounts using the cloud proxy. Each cloud proxy might have one or more cloud accounts. You can also view the health and status of these cloud accounts from this details pane. |

5. If your cloud proxy is not collecting data, you can view the health of the cloud proxy. From the left menu, click **Environment › Inventory**, select the **VMware Aria Operations Cloud Proxy Object** from the list, and then click **Show Detail**.

   For more details, see Inventory Tab and Inventory: List of Objects.

6. After you locate the VMware Aria Operations Cloud Proxy object, you can view the object details using the Summary tab. For more information, see Summary Tab.
7. Use the Alerts tab to monitor the health of the cloud proxy. If there are any issues, troubleshoot them using the Metrics tab.

   If your cloud proxy is not working properly, an alert is displayed.
   ```
   One or more VMware Aria Operations services on a cloud proxy are down
   ```

   To clear this alert, perform the following steps:
   - Check the network connectivity and configuration for the cloud proxy.
   - Take the cloud proxy offline and then bring it online.

   For more information, see Cloud Proxy FAQ and Cloud Proxy Troubleshooting.

   **NOTE**
   It is recommended that you create a notification rule for this alert so that, quick remediation steps can be taken, if necessary.

8. You can use the cloud proxy command line interface for other cloud proxy related actions. For more details, see Using the Cloud Proxy Command-Line Interface.

Click the **Collector Groups** tab to view the cloud proxies that are part of the collector groups. For more information on collector groups, see the "Managing Collector Groups" topic in the *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide*.

Click the **Collector Groups** tab to view the cloud proxies that are part of the collector groups. For more information on collector groups, see Managing Collector Groups.

### Deleting Cloud Proxies

You can delete a cloud proxy if there are issues with your cloud proxy. For example, you can delete the cloud proxy from the cloud proxy page in case of incorrect deployment of your cloud proxy, or if you need to deploy your cloud proxy in another cluster, or if you no longer require a cloud proxy.

1. Log in to VMware Cloud Foundation OperationsVMware Aria Operations.
2. From the left menu, click **Administration › Cloud Proxy**.
3. Click the vertical ellipsis in front of the cloud proxy you want to delete and then click **Delete**.

The confirmation wizard appears.

4. Select the **I understand the risk** check box and then click **OK**.
5. The cloud proxy is deleted.

### Upgrading Cloud Proxy

Cloud Proxies are upgraded to a compatible cluster version automatically after the cluster upgrade. Expect a downtime of one or two cycles, as cloud proxies do not collect any data during this period. Data collection resumes after the upgrade is complete. The upgrade process runs concurrently for up to ten cloud proxies. If you have more than ten cloud proxies, the upgrade process will continue sequentially and the next cloud proxy upgrade will start as soon as any one of the ten running upgrades completes or fails the upgrade process. In case the automatic upgrade fails, you can upgrade your cloud proxy manually using the CLI.

> **NOTE**
> When your cloud proxies are upgraded, the cloud proxy internal certificates also get renewed. After each upgrade, the cloud proxies will have a new root-CA certificate with a five year validity period.

For more information on what data gets collected, see Vmware vSphere Solution in VMware Aria Operations.

You can manually upgrade your cloud proxy Using the Cloud Proxy Command-Line Interface.

### Using the Cloud Proxy Command-Line Interface

You can use SSH to access the cloud proxy instance and use its Command-Line Interface to run the following actions:

- Manually upgrade your cloud proxy in case the automatic download of the latest binary fails. When the automatic download fails, you see a notification on the VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations user interface. To manually upgrade your cloud proxy instance to the latest version, see the following KB article 80590.
- Generate the support bundle.
- Gather the status of the cloud proxy's health and connectivity details.

| Command Line | Description |
|---|---|
| `cprc-cli -h, --help` | Displays the help message and use of the command-line interface. |
| `cprc-cli -s, --status` | Prints the cloud proxy life-cycle status, configuration details, upgrade-related information, and more. It is useful to catch necessary information related to support and troubleshooting, or check the connection to VMware Aria OperationsVMware Cloud Foundation OperationsVMware Aria Operations, or check the product version number, and so on. |
| `cprc-cli -u PRODUCT_PAK, --upgrade PRODUCT_PAK` | The cloud proxy instance is enabled for an automated upgrade by default. But if the automated upgrade fails due to any exceptional issue, use this command line to upgrade your cloud proxy instance to the desired version. |
| • 8.3 Release `cprc-cli -sb, --generate-support-bundle`<br>• 8.4 Release `cprc-cli -sb, --generate-support-bundle` | Generates the cloud proxy support bundle which is a package of logs, configurations, and status files. The support bundles are necessary for product support and troubleshooting. Generated support bundles can be found |

*Table continued on next page*

*Continued from previous page*

| Command Line | Description |
|---|---|
| • 8.5 Release `cprc-cli IS_HEAVY -sb, --generate-support-bundle IS_HEAVY`<br>The `IS_HEAVY` option should be specified as true or false. For example:<br>`cprc-cli -sb true`<br>`cprc-cli -sb false`<br>With the true option, the support bundle is generated with journal.ctl logs. With the false option, the support bundle is generated without journal.ctl logs | at the `/storage/core/vmware-vrops-cprc/support` location. |
| `cprc-cli -rsb SUPPORT_BUNDLE, --remove-support-bundle SUPPORT_BUNDLE` | Removes any specified support bundle. Although generated support bundle packages can be removed using the system-embedded commands, it is recommended to use this command for that action. |
| `cprc-cli -fm, --enable-fips-mode` | Enables FIPS mode for cloud proxy. |
| `cprc-cli -v -d all \| install \| network \| services \| connection` | Runs the troubleshoot process to check for all or any one of the following issues:<br>• Installation. Checks if VM deployment and cloud proxy installation is successful.<br>• Network. Checks IP configurations, DNS configurations, and Host (hostname) configurations.<br>• Service. Checks whether the `vmware-casa.service`, `collector.service`, `haproxy.service`, `httpd-north.service`, and `httpd-south.service` are active and loaded based on the `systemctl show` command.<br>• Connection: Checks the connection to VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations cluster nodes using `curl`. For the proxy server cases, checks that connection via the proxy server using its specified configuration.<br>Use the `-v` command to generate more verbose outputs. This command generates a log file that gets stored in `/var/log/support/cprc.analysis.log` location. |
| `-cs, --certificates-status` | Displays the status of the imported certificates and if the certificates have been configured correctly or not. |
| `cprc-cli -rc, --replace-certificate` | Allows you to replace the self signed certificates for cloud proxy.<br>Add the full path to the new certificate to the command. The function expects a certificate full chain containing the certificate and private key. For more information, see KB article 89583. |
| `cprc-cli -cc, --connectivity-check` | Allows you to run the connectivity check. You can check the HTTPS connectivity to a specified url. Incase the url is |

*Table continued on next page*

*Continued from previous page*

| Command Line | Description |
|---|---|
| | not specified, it checks the HTTPS connectivity to the default prerequisites<br>For on-prem configurations, the default prerequisites are the cluster nodes.<br><br>For SaaS configurations, the default prerequisites are the gateway FQDN and the upgrade pack url. |
| `cprc-cli -cc -r, --response` | Displays the response from the server during the connectivity check. |
| `cprc-cli -cc -rt, --routing` | Displays the routing information during the connectivity check. |
| `cprc-cli -env, --ovfenv` | Displays the content of ovfEnv.xml file. |
| `cprc-cli -rec, --recovery` | Runs the CPRC recovery script. |
| `cprc-cli -rh, --run-on-hosts` | Runs a specific script on hosts according to the configurations provided. |
| `cprc-cli -gt, --generate-template` | Generates the template.json files for -rh options configurations. |

## Using APIs with VMware Aria OperationsVMware Cloud Foundation Operations Cloud Proxy

You can use a browser or an HTTP client program to send requests and receive responses.

### REST Client Programs

Any client application that can send HTTPS requests is an appropriate tool for developing REST applications with the VMware Aria OperationsVMware Cloud Foundation Operations API. REST client plug-ins are available for most browsers and many IDEs. The following open-source programs are commonly used:

- cURL. http://curl.haxx.se
- Postman application. http://www.getpostman.com

In addition, VMware provides language-specific client bindings for the VMware Aria OperationsVMware Cloud Foundation Operations API.

### About the Schema Reference

The VMware Aria OperationsVMware Cloud Foundation Operations REST API documentation includes reference material for all elements, types, queries, and operations in the VMware Aria OperationsVMware Cloud Foundation Operations API. It also includes the schema definition files.

Swagger based API documentation is available with the product, with the capability of making REST API calls right from the landing page.

To access the API documentation, use the URL of your VMware Aria OperationsVMware Cloud Foundation Operations cloud proxy. For example, if the URL of your cloud proxy is `https://cloudproxy.vmware.comhttps://www.cloudproxy.vmware.com`, the API reference is available from: `https://cloudproxy.vmware.com/suite-api/doc/swagger-ui.htmlhttps://cloudproxy.vmware.com/suite-api/doc/swagger-ui.html`.

> **NOTE**
> You must use your cloud proxy IP or FDQN in the url to access the API documentation.

Language-specific client bindings are available from:

```
https://cloudproxy.vmware.com/suite-api/

https://www.cloudproxy.vmware.com/vrops-cloud/suite-api/
```

### Authorize VMware Aria OperationsVMware Cloud Foundation Operations API

To perform API calls, you must authorize your VMware Cloud Foundation Operations cloud proxy in the Swagger UI using the access token. Enter the **CSP Token (apiKey)** value and click **Authorize**.

To perform API calls, you must authorize your VMware Aria Operations cloud proxy. You can perform basic authorization and enter the **Username**, **Password**, and the **Auth Source**, or perform token based authorization and enter the **Ops Token (apikey)** value, and then click **Authorize**.

### About the VMware Aria OperationsVMware Cloud Foundation Operations API Examples

All examples include HTTP requests and responses. These examples show the workflow and content associated with operations such as creating and querying for information about objects in your monitored environment.

Example request bodies are in JSON. Request headers required by the VMware Aria OperationsVMware Cloud Foundation Operations API are included in example requests that are not fragments of a larger example.

Most example responses show only those elements and attributes that are relevant to the operation being discussed. Ellipses (...) indicate omitted content within response bodies.

## Cloud Proxy FAQ

This topic covers some frequently asked questions about VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations Proxy.

### Configuration

1. What are the prerequisites for setting up a cloud proxy account?
   Prerequisites are given in the topic, Configuring Cloud Proxies in .
2. What does one-way connection mean?
   Outbound connections are initiated from cloud proxy to VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations, over **`https/443`**. Cloud proxy can also facilitate vCenter actions.
3. Which ports should be opened?
   The most up-to-date technical information about ports can be found on Ports and Protocol.
4. Which ports should be allowed for incoming traffic to cloud proxy?
   Allow port 443 https protocol for push model adapters like application monitoring or Suite-API on cloud proxy. Allow ports 4505, 4506, and 8443 via TCP protocol for application monitoring. Allow the VRRP protocol for intercommunication between cloud proxies in a application monitoring high availability activated collector group.
5. How do I edit environment settings for cloud proxy?
   You can edit vApp options. For more information, see Edit OVF Details for a Virtual Machine.
6. How are certificates managed?
   Certificates are managed by cloud proxies. But for any additional proxy servers with SSL communication, you need to provide certificate(s).
7. What credential is used to login to cloud proxy?
   You can login as the “root” user. You are expected to set a new password on the first login to cloud proxy VM.

   SSH access is disabled by default, so the first login must be done via the vCenter console. You can run the following command to start SSH service:
   ```
   systemctl start sshd

   systemctl enable sshd
   ```

To reset password, see the VMware KB Article, 2001476.

8. Where can I configure the local HTTP proxy for VMC on AWS?
   Perform the following steps:

   1. Login to VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations and go to the Administration page.
   2. Go to Cloud Accounts.
   3. Select VMC on AWS.
   4. Click + next to credentials to add a credential.
   5. In proxy details, add details for the local HTTP proxy. (Do not add details for cloud proxy here).

   For more details, see the Configuring VMware Cloud on AWS in VMware Aria Operations topic in the *VMware Aria Operations Configuration Guide*.
9. Will I be notified if the connection between cloud proxy and VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations breaks down?
   You can configure alerts/notifications on the *VMware Aria Operations cloud proxy* object. For more information, see Monitoring the Health of Cloud Proxies.

   VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations automatically generates notifications for the following scenarios:
   - Cloud proxy is not reachable.
   - Cloud proxy is nearing sizing limits.
10. How do I change account for cloud proxy?
    You can edit vApp options. For more information, see Edit OVF Details for a Virtual Machine.
11. How can I check the status of cloud proxy?
    For more information, see Monitoring the Health of Cloud Proxies.
12. Should I use Remote Collector or cloud proxy for monitoring?
    VMware recommends that you use cloud proxy to take advantage of the latest enhancements. Also, application monitoring, HA of collector groups, and data persistence are only supported through cloud proxy.

## Sizing

1. How should I size the cloud proxy?
   For information on sizing, see the VMware KB article 78491

   For information on sizing, see the VMware KB article 85832
2. How would I know if cloud proxy is nearing sizing limit?
   VMware Cloud Foundation OperationsVMware Aria OperationsVMware Aria Operations customers will receive an email when cloud proxy is nearing sizing limit.

## Upgrade

1. How do I upgrade cloud proxy?
   Cloud proxy is upgraded automatically. In case the upgrade fails, see the VMware KB article 80590.

## High Availability

1. Is high availability supported?
   Cloud proxy supports high availability. You can add multiple cloud proxies to a collector group. If the collecting cloud proxy fails or gets disconnected, collection can be picked up by another proxy in the group.

   **NOTE**
   Since the failover is initiated after a period of 10 minutes, few collection cycles are lost.

To troubleshoot cloud proxy issues, see Cloud Proxy Troubleshooting.

## Cloud Proxy Troubleshooting

Cloud proxy troubleshooting steps are provided to help you easily resolve issues that you may come across in VMware Aria Operations.

Before you proceed with troubleshooting, see the Cloud Proxy FAQ.

### Installation and/or First Boot Failure

`To verify the issue, check if/var/log/firstbootcontains a file named "Succeeded".`

If not, the following problems could result in VMware Aria Operationsinstallation and/or first boot failure:

1. Unique Registration Key used while deploying Cloud Proxy is invalid. To verify, check the cloud proxy console. Solution: Redeploy cloud proxy.

### Cloud Proxy VM is running, but the status is Offline in VMware Aria Operations.

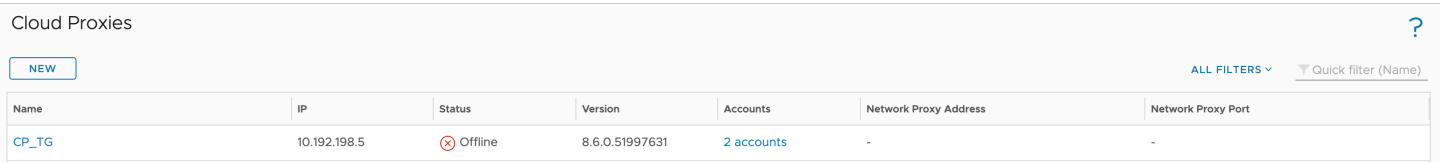


To verify the connection, use the following commands: (For the complete list of commands, please see Using the Cloud Proxy Command-Line Interface.

```
# Overall status of cloud proxy:cprc-cli -s


# Ping itself:

ip addr

ping <address>



# Ping gateway:

ip route

ping <gateway>


# Verify the connection outside the cloud proxy,

ping 8.8.8.8


Note: If you are using a network proxy,

use the /opt/vmware/share/vami/vami_config_net option#5 command
```

```
to ensure you have the correct configuration for the testings.
```

The following problems could result in VMware Aria Operationsdisplaying the status of cloud proxy as offline.

1. Incorrect network proxy information in cloud proxy configuration.
   To verify the connection via a network proxy, use the following:
   ```
   curl -vvv --proxy http(s)://proxy_user:proxy_pass@proxy_ip:proxy_port -H 'Accept:
   application/json' -H 'Content-Type: application/json' -X GET https://<gateway url>/
   casa/security/ping (gateway url example - 10238.gw.dev.vrops-ops.com)
   ```

   ```
   To ignore SSL validation for a proxy server,

   use curl --proxy-insecure. With SSL validation the customer can provide Proxy Server
   certificate during cloud proxy deployment or re-configuration

   so that provided certificate from customer can be used to check the connection with
   curl with SSL certificate validation.
   ```

   Solution:

   1. SSH to the Cloud Proxy VM and set the `connectretry` to `0` in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.
   2. Shutdown the Cloud Proxy VM.
   3. Update the network proxy configurations from the vCenter VM options using the vApp options Edit OVF Details for a Virtual Machine.
   4. Boot the Cloud Proxy VM.

2. Required ports are not open.
   To verify:
   ```
   openssl s_client -showcerts -connect {address}:443
   ```

   ```
   curl -v telnet://{address}:443
   ```

   ```
   # Or, change the address to the machine you want to check:

   python -c "import socket; print(socket.socket(socket.AF_INET,
   socket.SOCK_STREAM).connect_ex(('127.0.0.1', 443)))"
   ```

   ```
   # If you get a !=0 response, the server is not listening to the port.
   ```

   Solution:

   1. SSH to the Cloud Proxy VM and set `connectretry` to `0` in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.
   2. Provide port access as mentioned in the prerequisite section of Configuring Cloud Proxies in
   3. Boot the Cloud Proxy VM.

3. Invalid certificate.
   To verify:

```
openssl s_client -showcerts -connect {address}:443
```

Solution:

1. SSH to the Cloud Proxy VM and set `connectretry` to `0` in `/storage/db/vmware-vrops-cprc/configuration/cprc.configuration` to ensure that the Cloud Proxy retries to connect.
2. Follow the steps mentioned in VMware KB Article, 83698.

4. The logs folder **`/storage/log`** is running out of partition space.
   Solution: Remove log files to ensure that enough space is available. Note that this is an exceptional case. In normal conditions, log files are auto archived.
5. One or more of the following services are down: `httpd-north.service`, `haproxy.service` and `collector.service`.
   Solution:
   - Check service status by running the following command: **`systemctl status <service name>`**.
   - To start service, use the following command: **`systemctl start <service name>`**.
6. Unique Registration Key expired.
   Solution: Redeploy Cloud Proxy with new Unique Registration Key.

### Cloud proxy is online, and state of Cloud Account is `Collecting`, but status is `Object Down`.

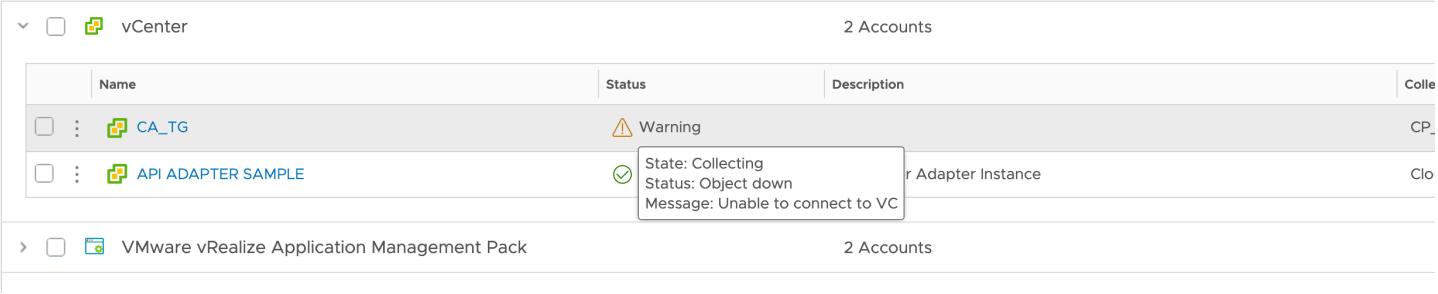


The following problem could result in VMware Aria Operationsdisplaying the state of Cloud Account as `Collecting`, while the status is, `Object Down`.

1. Incorrect account credentials.
   Solution: Check and update the credentials used while setting up the cloud account.

### Cloud proxy status is stuck in `Going Online`.

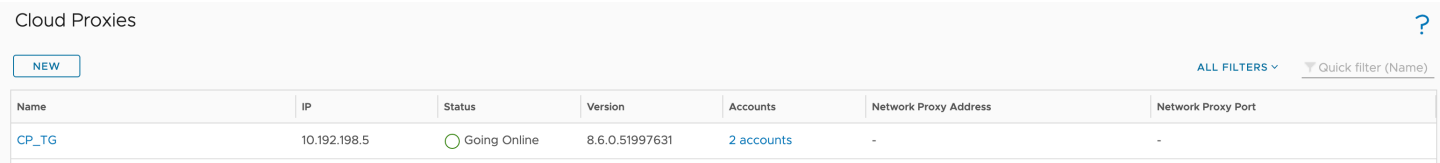


It can take up to 20 mins on first reboot, for the cloud proxy to be registered and come online. Wait for the specified time to see if cloud proxy comes online. If it still does not come online, one or more of the following services are down: `httpd-north.service`, `haproxy.service`, and `collector.service`.

Solution:

1. Check service status by running the following command: `systemctl status <service name>`
2. To start service, use the following command: `systemctl start <service name>`.

**Cloud proxy does not upgrade automatically, after the upgrade of VMware Aria Operations**

There could be a few possible reasons why cloud proxy does not upgrade automatically after an upgrade of VMware Aria Operations.

1. High network latency leading to PAK download failure. Latency of >500ms is not supported.
   Solution: See the VMWare KB article 80590 on how to manually upgrade cloud proxy via CLI.
2. Upgrade status is stuck at `Running` since the previous upgrade had failed.
   Solution: Follow the steps given below to change the upgrade status.

   1. Stop the casa service: **`systemctl stop vmware-casa.service`**.
   2. Change the upgrade status from `RUNNING` to `NONE` in the following files:
      `./storage/db/vmware-vrops-cprc/status/cprc.upgrade.status`

      `./storage/db/vmware-vrops-cprc/status/cprc.pak.status`
   3. See the VMware KB article 80590 and run the manual upgrade.

**Cloud proxy gets disconnected at regular intervals**

There could be a few possible reasons why cloud proxy gets disconnected at regular intervals.

1. Check the network connectivity and latency.
2. Check if the cloud proxy VM can reach the DNS and use the NSlookup to validate the DNS connectivity.

## Configuring Collector Groups

VMware Aria OperationsVMware Cloud Foundation Operations uses collectors like cloud proxies to manage cloud account processes such as gathering metrics from objects. You can select a collector or a collector group when configuring an cloud account.

If there are cloud proxies in your environment, you can create a collector group, and add thecloud proxies to the group. When you assign an account to a collector group, that account can use any collector in the group. Use collector groups to achieve resiliency in cases where the collector experiences network interruption or becomes unavailable. If this occurs, and the collector is part of a group, the total workload is redistributed among all the collectors in the group, reducing the workload on each collector. For more information, see 'Configuring Cloud Proxies in VMware Aria Operations' topic in the *Configuring VMware Aria Operations Guide*.

If there are cloud proxies in your environment, you can create a collector group, and add cloud proxies to the group. When you assign an account to a collector group, the account can use any collector in the group. Use collector groups to achieve high availability in cases where the collector experiences network interruption or becomes unavailable. If this occurs, and the collector is part of a group, the total workload is redistributed among all the collectors in the group, reducing the workload on each collector. For more information, see Configuring Cloud Proxies in VMware Aria Operations .

You can add, edit, or remove collector groups in VMware Aria OperationsVMware Cloud Foundation Operations, and rebalance your cloud accounts.

**Rebalancing a Cloud Account**

Rebalancing of your cloud accounts is not intended to provide equally distributed cloud accounts across each collector in the collector group. The rebalancing action considers the number of resources that each cloud account collects to determine the rebalancing placement. The rebalancing happens at the cloud account, which can result in several small cloud accounts on a single collector, and a single huge cloud account on another collector, in your VMware Cloud Foundation OperationsVMware Aria Operations instance.

Rebalancing your collector groups can add a significant load on the entire cluster. Moving cloud accounts from one collector to another collector requires that VMware Cloud Foundation OperationsVMware Aria Operations stops the cloud account and all its resources on the source collector, then starts them on the target collector.

If a collector fails to respond or loses connectivity to the cluster, VMware Cloud Foundation OperationsVMware Aria Operations starts automated rebalancing in the collector group. All other user-initiated manual operations on the collector, such as stopping or restarting the collector manually, do not result in automated rebalancing.

If one of the collectors fails to respond, or if it loses network connectivity, VMware Cloud Foundation OperationsVMware Aria Operations performs automated rebalancing. In cases of automated rebalancing, to properly rebalance the collector group, you must have spare capacity on the collectors in the collector group.

**Where You Manage Collector Groups**

From the left menu, click **Administration › Cloud Proxies**, and then click the **Collector Groups** tab.

**Table 25: Collector Group Summary Grid**

| Options | Description |
|---|---|
| Collector Group toolbar | To manage collector groups, use the toolbar icons.<br>• **Add**. Add a collector group<br>• Click the **Vertical Ellipses** to perform any one of the following actions:<br>– **Edit**. Modify the collector group by adding or removing cloud proxies. You can also edit a collector group to activate or deactivate application monitoring high availability.<br>– **Delete**. Remove the selected collector group.<br>– **Rebalance collector group**. Rebalance one collector group at a time. If you have permission to manage clusters, you can rebalance the workload across the collectors and the cloud proxies in the collector group. The rebalance action moves objects from one collector group to another to rebalance the number of objects on each collector in the collector group. If a disk rebalance is already in progress, the collector rebalance does not run<br>– **Retry collector group configuration**: In case the application monitoring high availability activated collector group configuration fails, you can retry the configuration.<br>– The actions are disabled when the application monitoring high availability activated collector group configuration is in progress. Once the cofiguration is complete you can perform the actions. |
| Name | The name given to the collector group when the collector group is created or updated. |
| Collector Group ID | The collector group ID created when the collector group is created. |
| Description | Description given to the collector group when the collector group is created or updated. |
| Application Monitoring HA | Displays the application monitoring high availability status of the collector group.<br>• Activated. The collector group can be used for application monitoring.<br>• Deactivated. The collector group cannot be used for application monitoring.<br>**NOTE**<br>For default collector groups, high availability for application monitoring is deactivated by default. |
| HA Status | Displays the application monitoring high availability of the collector group.<br>• OK. The collector group high availability configuration is successful. |

*Table continued on next page*

*Continued from previous page*

| Options | Description |
|---|---|
| | • In Progress. The collector group high availability configuration is in progress. This status is displayed when adding, deleting, updating, or deactivating the collector group.<br>• Empty. If high availability is not configured, no status is displayed.<br><br>**NOTE**<br>If the high availability status of the cloud proxy fails, the HA status of the failed cloud proxy is displayed. In case there are multiple cloud proxies in a collector group with failed HA status, the first failed message is displayed in the Collectors Group page. |
| Virtual IP | Virtual IP address of the collector group.<br><br>**NOTE**<br>Pick the Virtual IP address that any of the cloud proxies in a collector group can own and receive traffic from. The Virtual IP address must be in the same subnet as the physical address of the cloud proxies. After you configure the Virtual IP address, ping it from a different network to ensure that it is routable. |
| Filters | Search the list of collector groups according to the following criteria:<br>• Name<br>• Description<br>• Collector Name<br>• Collector IP Address<br>• Application Monitoring HA<br>• HA Status<br>• Group Virtual IP<br>• Collector Group Id |

Click a collector group to view the collector group details.

**Table 26: Collector Group Details Grid**

| Detail Grid Options | Description |
|---|---|
| Collector Group Name | Displays the name of the collector group. Click **Edit** to modify the collector group details. Click **View Description** to view the collector group description. |
| Application Monitoring High Availability Status | Displays the status of the application monitoring high availability. |
| Virtual IP | Displays the virtual IP added during the collector group creation. |
| Name | Name given to the cloud proxy when the cloud proxy was created. Click the name to view the cloud proxy details. . For more information on cloud proxies, see the "Monitoring the Health of Cloud Proxies" topic in the *Configuring VMware Aria OperationsVMware Cloud Foundation Operations Guide*. |

*Table continued on next page*

*Continued from previous page*

| Detail Grid Options | Description |
|---|---|
| | For more information on cloud proxies, see Monitoring the Health of Cloud Proxies. |
| IP Address | IP address of the cloud proxy |
| Status | Status of the cloud proxy: online or offline. |
| HA Status | Displays the cloud proxy high availability status.<br>• OK. The cloud proxy application monitoring high availability configuration is successful.<br>• Activation In Progress. Displayed when the cloud proxy is added to a application monitoring HA activated collector group or when the creation of a new application monitoring HA activated collector group is in progress.<br>• Deactivation In Progress. Displayed when the deactivation of a application monitoring HA activated collector group is in progress.<br>• Removal In Progress. Displayed when the removal of a cloud proxy from a application monitoring HA activated collector group is in progress.<br>• Group Delete In Progress. Displayed when the deletion of the entire application monitoring HA activated collector group is in progress.<br>• Empty. If the cloud proxy is not part of a collector group with high availability, no status is displayed.<br>• If the cloud proxy high availability configuration fails, the failure reason is displayed. For example, "Keepalived is failed" or "Apache south is failed". |
| Type | Displays the cloud proxy type. |
| Number of Objects | Displays the number of objects collected by the cloud proxy. |
| Number of cloud accounts | Displays the number of cloud accounts using the cloud proxy. |

### Adding a Collector Group

Create a new collector group from the available cloud proxies in your VMware Cloud Foundation Operations environment. A cloud proxy can only be added to only one collector group at a time.

**Where You Add New Collector Groups?**

From the left menu, click **Administration › Cloud Proxies**, and then click the **Collector Groups** tab. Click **Add**.

**Table 27: Add New Collector Group**

| Option | Description |
|---|---|
| Name | Name of the collector group. |
| Description | Description of the collector group. |
| Application Monitoring High Availability | Activate this option to use high availability in application monitoring using collector groups. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Virtual IP | If you activate high availability for application monitoring, enter the virtual IP address.<br>**NOTE**<br>Pick the Virtual IP address that any of the cloud proxies in a collector group can own and receive traffic from. The Virtual IP address must be in the same subnet as the physical address of the cloud proxies. After you configure the Virtual IP address, ping it from a different network to ensure that it is routable. |
| Collectors added to this collector group | Displays the collectors that are assigned to this collector group.<br>To assign collectors, double click a collector, or drag and drop it from the **Select Collectors** section to add it to the collector group.<br>**NOTE**<br>You cannot assign a collector that is part of another collector group. To reassign a collector to another collector group, you must remove it from the existing collector group and then reassign it.<br>To remove a collector from the collector group, click the **Remove Collector** icon. |
| Select Collectors | Displays a list of cloud proxies in your VMware Cloud Foundation Operations environment together with their collector name, IP address, type, collector group, and status.<br>Collectors that are assigned to this collector group appear with a check mark before the collector name.<br>**NOTE**<br>Cloud proxies with log forwarding cannot be added to collector groups. |
| Filters | You can search the list of collectors according to the following criteria:<br>• Collector Name<br>• IP address<br>• Collector Group Name<br>• Status |

### Editing Collector Groups

Edit a collector group in VMware Cloud Foundation Operations by adding cloud proxies to the group, or removing the collectors that you no longer require to be part of the group. You can also edit a collector group to activate or deactivate application monitoring high availability.Edit a collector group in VMware Aria Operations by adding cloud proxies to the group, or removing the collectors that you no longer require to be part of the group.

**Where You Edit a Collector Group?**

From the left menu, click **Administration › Cloud Proxies**, and then click the **Collector Groups** tab. Click the vertical ellipsis and then select **Edit**. The Edit Collector Group page opens. You can update the collector group details and then click **Save**.

**Table 28: Edit Collector Group Options**

| Option | Description |
|---|---|
| Name | Name given to the collector group when the collector group is created. |
| Description | Description given to the collector group when the collector group is created. |
| Application Monitoring High Availability | Activate this option to use high availability in application monitoring using collector groups.<br>If high availability in application monitoring is activated, you can deactivate it. |
| Virtual IP | If you activate high availability for application monitoring, enter the virtual IP.<br>**NOTE**<br>You cannot edit the virtual IP address for a collector group with high availability for application monitoring activated. |
| Collectors added to this collector group | Displays the collectors that are assigned to this collector group.<br>To assign collectors, double click a collector, or drag and drop it from the **Select Collector** section to add it to the collector group.<br>**NOTE**<br>You cannot assign a collector that is part of another collector group. To reassign a collector to another collector group, you must remove it from the existing collector group and then reassign it.<br>To remove a collector from the collector group, click the **Remove Collector** icon. |
| Select Collectors | Displays a list of the available cloud proxies in your VMware Cloud Foundation Operations environment together with their name, IP address, type, collector group, and status.<br>Collectors that are assigned to this collector group appear with a check mark before the collector name.<br>**NOTE**<br>Cloud proxies with log forwarding cannot be added to collector groups. |
| All Filters | Filter the list of collectors according to the following criteria:<br>• Collector Name |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Collector Group Name<br>• IP Address<br>• Status |

## Configuring Policies

To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more groups of objects.

## Policies

A policy is a set of rules that you define for VMware Aria OperationsVMware Cloud Foundation Operations to use to analyze and display information about the objects in your environment. You can create, modify, and administer policies to determine how VMware Aria OperationsVMware Cloud Foundation Operations displays data in dashboards, views, and reports.

### How Policies Relate to Your Environment

VMware Aria OperationsVMware Cloud Foundation Operations policies support the operational decisions established for your IT infrastructure and business units. With policies, you control what data VMware Aria OperationsVMware Cloud Foundation Operations collects and reports on for specific objects in your environment. Each policy can inherit settings from other policies, and you can customize and override various analysis settings, alert definitions, and symptom definitions for specific object types, to support the service Level agreements and business priorities established for your environment.

When you manage policies, you must understand the operational priorities for your environment, and the tolerances for alerts and symptoms to meet the requirements for your business critical applications. Then, you can configure the policies so that you apply the correct policy and threshold settings for your production and test environments.

Policies define the settings that VMware Aria OperationsVMware Cloud Foundation Operations applies to your objects when it collects data from your environment. VMware Aria OperationsVMware Cloud Foundation Operations applies policies to newly discovered objects, such as the objects in an object group. For example, you have an existing VMware adapter instance, and you apply a specific policy to the group named World. When a user adds a new virtual machine to the vCenter instance, the VMware adapter reports the virtual machine object to VMware Aria OperationsVMware Cloud Foundation Operations. The VMware adapter applies the same policy to that object, because it is a member of the World object group.

To implement capacity policy settings, you must understand the requirements and tolerances for your environment, such as CPU use. Then, you can configure your object groups and policies according to your environment.

- For a production environment policy, a good practice is to configure higher performance settings, and to account for peak use times.
- For a test environment policy, a good practice is to configure higher utilization settings.

VMware Aria OperationsVMware Cloud Foundation Operations applies the policies in the priority order, as they appear in the priority column. When you establish the priority for your policies, VMware Aria OperationsVMware Cloud Foundation Operations applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of any active policy:

1. In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

**NOTE**
The Reorder Policies option is activated only if there are more than one active policies.

2. In the Reorder Policies window, select the policy and drag it up or down to change the priority.
3. Click **ok** to save the changes made to the priority.

The priority for the Default Policy is always designated with the letter D, and the other active policies are prioritized with numbers 1, 2, and so on. Policy with priority 1 indicates the highest priority. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, VMware Aria OperationsVMware Cloud Foundation Operations associates the highest ranking policy with that object.

**Table 29: Configurable Policy Rule Elements**

| Policy Rule Elements | Thresholds, Settings, Definitions |
|---|---|
| Workload | Configure symptom thresholds for Workload. |
| Time Remaining | Configure thresholds for the Time Remaining. |
| Capacity Remaining | Configure thresholds for the Capacity Remaining. |
| Maintenance Schedule | Sets a time to perform maintenance tasks. |
| Attributes | An attribute is a collectible data component. You can activate or deactivate metric, property, and super metric attributes for collection, and set attributes as key performance indicators (KPIs). A KPI is the designation of an attribute that indicates that the attribute is important in your own environment. |
| Alert Definitions | Activate or deactivate combinations of symptoms and recommendations to identify a condition that classifies as a problem. |
| Symptom Definitions | Activate or deactivate test conditions on properties, metrics, or events. |

**Privileges to Create, Modify, and Prioritize Policies**

You must have privileges to access specific features in the VMware Aria OperationsVMware Cloud Foundation Operations user interface. The roles associated with your user account determine the features you can access and the actions you can perform. To set the policy priority:

1. In the Policies page, click the horizontal ellipse, and click **Reorder Policies**.

   **NOTE**
   The Reorder Policies option is activated only if there are more than one active policies.

2. In the Reorder Policies window, select the policy and drag it up or down to change the priority.
3. Click **ok** to save the changes made to the priority.

**How Upgrades Affect Your Policies**

After you upgrade VMware Aria OperationsVMware Cloud Foundation Operations from a previous version, you might find newly added or updated default settings of policies such as, new alerts and symptoms. Hence, you must analyze the settings and modify these settings to optimize them for your current environment. If you apply the policies used with a previous version of VMware Aria OperationsVMware Cloud Foundation Operations, the manually modified policy settings remain unaltered.

### Policy Decisions and Objectives

Implementing policy decisions in VMware Aria OperationsVMware Cloud Foundation Operations is typically the responsibility of the Infrastructure Administrator or the Virtual Infrastructure Administrator, but users who have privileges can also create and modify policies.

You must be aware of the policies established to analyze and monitor the resources in your IT infrastructure.

- If you are a Network Operations engineer, you must understand how policies affect the data that VMware Aria OperationsVMware Cloud Foundation Operations reports on objects, and which policies assigned to objects report alerts and issues.
- If you are the person whose role is to recommend an initial setup for policies, you typically edit and configure the policies in VMware Aria OperationsVMware Cloud Foundation Operations.
- If your primary role is to assess problems that occur in your environment, but you do not have the responsibility to change the policies, you must still understand how the policies applied to objects affect the data that appears in VMware Aria OperationsVMware Cloud Foundation Operations. For example, you might need to know which policies apply to objects that are associated with particular alerts.
- If you are a typical application user who receives reports from VMware Aria OperationsVMware Cloud Foundation Operations, you must have a high-level understanding of the operational policies so that you can understand the reported data values.

**Policies Library**

The policies library displays the base settings, default policy, and other best practice policies that VMware Aria OperationsVMware Cloud Foundation Operations includes. You can use the policies library to create your own policies. The policies library includes all the configurable settings for the policy elements, such as workload, capacity and time remaining, and so on.

**How the Policies Library Works**

Use the options in policies library to create your own policy from an existing policy, or to override the settings from an existing policy so that you can apply the new settings to groups of objects. You can also import or export a policy and reorder the policies.

Select a policy to display its details in the right pane. The right pane displays a high-level overview of all the details and options for that policy where these details are categorized in tabs. Expand each category to view all the related details.

When you add or edit a policy, you access the policy workspace where you select the base policies and override the settings for metrics and properties, alerts and symptoms, capacity, compliance, workload automation, and groups and objects. In this workspace, you can also apply the policy to objects and object groups. To update the policy associated with an object or object group, the role assigned to your user account must have the Manage Association permission activated for policy management.

**Where You Manage the Policies Library**

To manage the policies library, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. The policies library appears and lists the policies available to use for your environment.

**Table 30: Policy Library Tab Options**

| Option | Description |
|---|---|
| Toolbar | Use the toolbar selections to take action in the policies library.<br>• Add. Create a policy from an existing policy.<br>• Edit. Customize the policy so that you can override settings for VMware Aria OperationsVMware Cloud Foundation Operations to analyze and report data about the associated objects.<br>• Delete. Remove a policy from the list.<br>• Set Default Policy. You can set any policy to be the default policy, which applies the settings in that policy to all objects that do not have a policy applied. When you set a policy to be the default policy, the priority is set to `D`, which gives that policy the highest priority. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Export. Downloads the policy.<br>• Import. Allows you to import policies. To import:<br>– Click the **Import** option from the horizontal ellipsis.<br>– Click **Browse** and select the file to import.<br>– Select if you want to Overwrite or Skip the file in case of a conflict.<br>– Click **Import** to import the policy, and click **Done**.<br>**NOTE**<br>To import or export a policy, the role assigned to your user account must have the Import or Export permissions activated for policy management.<br>• Reorder Policies. Change the priority of the active policies. |
| Filters | Limits the list based on the text you type.<br>You can also filter by:<br>• Name<br>• Description<br>• Modified By |
| Policies library data grid | VMware Aria OperationsVMware Cloud Foundation Operations displays the high-level details for the policies.<br>• Name. Name of the policy as it appears in the Add or Edit Policy workspace, and in areas where the policy applies to objects, such as in Custom Groups.<br>• Status: Indicates whether the policy is active or inactive.<br>• Description. Meaningful description of the policy, such as which policy is inherited, and any specific information users need to understand the relationship of the policy to one or more groups of objects.<br>• Last Modified. Date and time that the policy was last modified. |
| Policies library > Right Pane | The right pane displays the name and description of the policy from which the settings are inherited, the policy priority, and the option to edit the policy. From the right pane, you can view the complete group of settings that include both customized settings and the settings inherited from the base policies selected when the policy was created.<br>• Metrics and Properties: Displays all the attribute types included in the policy. Attribute type includes, metrics properties, and super metrics.<br>• Alerts and Symptoms: Displays all the alert and symptom definitions included in the policy. The Alert Definitions tabs display an overview of the alert definition, criticality, symptom, and state. The Symptoms Definitions tab displays an overview of the symptom name, criticality, and the metric name.<br>• Capacity: Displays an overview of all the thresholds of the objects included in the policy.<br>• Compliance: Displays the compliance thresholds inherited from the base policy or set while creating the policy.<br>• Workload Automation: Displays the details of the workload optimized in your environment per your definition.<br>• Groups and Objects: Displays the object or object groups associated with the selected policy and the names of the objects in your environment, their object types, and associated adapters. When a parent group exists for an object, it is shown here. |

## Operational Policies

Determine how to have VMware Aria OperationsVMware Cloud Foundation Operations monitor your objects, and how to notify you about problems that occur with those objects.

VMware Aria OperationsVMware Cloud Foundation Operations Administrators assign policies to objects or object groups and applications to support Service Level Agreements (SLAs) and business priorities. When you use policies with objects or object groups, you ensure that the rules defined in the policies are quickly put into effect for the objects in your environment.

With policies, you can:

- Activate and deactivate alerts.
- Control data collections by persisting or not persisting metrics on the objects in your environment.
- Configure the product analytics and thresholds.
- Monitor objects and applications at different service levels.
- Prioritize policies so that the most important rules override the defaults.
- Understand the rules that affect the analytics.
- Understand which policies apply to objects or object groups.

VMware Aria OperationsVMware Cloud Foundation Operations includes a library of built-in active policies that are already defined for your use. VMware Aria OperationsVMware Cloud Foundation Operations applies these policies in priority order.

When you apply a policy to an object or an object group, VMware Aria OperationsVMware Cloud Foundation Operations collects data from the objects based on the thresholds, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that are activated in the policy.

The following examples of policies might exist for a typical IT environment.
- Maintenance: Optimized for ongoing monitoring, with no thresholds or alerts.
- Critical Production: Production environment ready, optimized for performance with sensitive alerting.
- Important Production: Production environment ready, optimized for performance with medium alerting.
- Batch Workloads: Optimized to process jobs.
- Test, Staging, and QA: Less critical settings, fewer alerts.
- Development: Less critical settings, no alerts.
- Low Priority: Ensures efficient use of resources.
- Default Policy: Default system settings.

## Types of Policies

There are three types of policies such as default policies, custom policies, and policies that are offered with VMware Aria OperationsVMware Cloud Foundation Operations.

### Custom Policies

You can customize the default policy and base policies included with VMware Aria OperationsVMware Cloud Foundation Operations for your own environment. You can then apply your custom policy to an individual object or groups of objects, such as the objects in a cluster, or virtual machines and hosts, or to a group that you create to include unique objects and specific criteria.

You must be familiar with the policies so that you can understand the data that appears in the user interface, because policies drive the results that appear in the VMware Aria OperationsVMware Cloud Foundation Operations dashboards, views, and reports.

To determine how to customize operational policies and apply them to your environment, you must plan ahead. For example:

- Must you track CPU allocation? If you overallocate CPU, what percentage must you apply to your production and test objects?
- Will you overallocate memory or storage? If you use High Availability, what buffers must you use?
- How do you classify your logically defined workloads, such as production clusters, test or development clusters, and clusters used for batch workloads? Or, do you include all clusters in a single workload?
- How do you capture peak use times or spikes in system activity? In some cases, you might need to reduce alerts so that they are meaningful when you apply policies.

When you have privileges applied to your user account through the roles assigned, you can create and modify policies, and apply them to objects. For example:

- Create a policy from an existing base policy, inherit the base policy settings, then override specific settings to analyze and monitor your objects.
- Use policies to analyze and monitor vCenter objects and non- vCenter objects.
- Set custom thresholds for capacity settings on all object types to have VMware Aria OperationsVMware Cloud Foundation Operations report on workload, and so on.
- Activate specific attributes for collection, including metrics, properties, and super metrics.
- Activate or deactivate alert definitions and symptom definitions in your custom policy settings.
- Apply the custom policy to an individual object or groups of objects.

When you use an existing policy to create a custom policy, you override the policy settings to meet your own needs. You set the allocation and demand, the overcommit ratios for CPU and memory, and the thresholds for capacity risk and buffers. To allocate and configure what your environment is actually using, you use the allocation model and the demand model together. Depending on the type of environment you monitor, such as a production environment versus a test or development environment, whether you over allocate at all and by how much depends on the workloads and environment to which the policy applies. You might be more conservative with the level of allocation in your test environment and less conservative in your production environment.

When you establish the priority for your policies, VMware Aria OperationsVMware Cloud Foundation Operations applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, VMware Aria OperationsVMware Cloud Foundation Operations associates the highest ranking policy with that object.

Your policies are unique to your environment. Because policies direct VMware Aria OperationsVMware Cloud Foundation Operations to monitor the objects in your environment, they are read-only and do not alter the state of your objects. For this reason, you can override the policy settings to fine-tune them until VMware Aria OperationsVMware Cloud Foundation Operations displays the results that are meaningful and that affect for your environment. For example, you can adjust the capacity buffer settings in your policy, and then view the data that appears in the dashboards to see the effect of the policy settings.

**Default Policy in VMware Aria OperationsVMware Cloud Foundation Operations**

The default policy is a set of rules that applies to most of your objects.

The Default policy is marked with the letter D in the Priority column and can apply to any number of objects.

All the Default policies appear in the Default Policy group in the policies library, even if that policy is not associated with an object group. When an object group does not have a policy applied, VMware Aria OperationsVMware Cloud Foundation Operations associates the Default policy with that group.

A policy can inherit the Default policy settings, and those settings can apply to various objects under several conditions.

The policy that is set to Default always takes the lowest priority. If you attempt to set two policies as the Default policy, the first policy that you set to Default is initially set to the lowest priority. When you set the second policy to Default, that policy then takes the lowest priority, and the earlier policy that you set to Default is set to the second lowest priority.

You can use the Default policy as the base policy to create your own custom policy. You modify the default policy settings to create a policy that meets your analysis and monitoring needs. When you start with the Default policy, your new policy inherits all the settings from the Default base policy. You can then customize your new policy and override these settings.

The data adapters and solutions installed in VMware Aria OperationsVMware Cloud Foundation Operations provide a collective group of base settings that apply to all objects. In the policy navigation tree in the policies library, these settings are called Base Settings. The Default policy inherits all the base settings by default.

**Policies Provided with VMware Aria OperationsVMware Cloud Foundation Operations**

VMware Aria OperationsVMware Cloud Foundation Operations includes sets of policies that you can use to monitor your environment, or as the starting point to create your own policies.

Verify that you are familiar with the policies provided with VMware Aria OperationsVMware Cloud Foundation Operations so that you can use them in your own environment, and to include settings in new policies that you create.

**Where You Find the Policies Provided with VMware Aria OperationsVMware Cloud Foundation Operations Policies**

From the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile to see the policies provided with VMware Aria OperationsVMware Cloud Foundation Operations.

**Policies That VMware Aria OperationsVMware Cloud Foundation Operations Includes**

All policies exist under the Base Settings, because the data adapters and solutions installed in your VMware Aria OperationsVMware Cloud Foundation Operations instance provide a collective group of base settings that apply to all objects. In the policies library, these settings are called Base Settings.

The Base Settings policy is the umbrella policy for all other policies, and appears at the top of the policy list in the policies library. All the other policies reside under the Base Settings, because the data adapters and solutions installed in your VMware Aria OperationsVMware Cloud Foundation Operations instance provide a collective group of base settings that apply to all objects.

The configuration based policy set includes policies provided with VMware Aria OperationsVMware Cloud Foundation Operations that you use for specific settings on objects to report on your objects. This set includes several types of policies:

- Efficiency alerts policies for infrastructure objects and virtual machines
- Health alerts policies for infrastructure objects
- Overcommit policies for CPU and Memory
- Risk alerts policies for infrastructure objects and virtual machines

The Default Policy includes a set of rules that applies to most of your objects.

## Using the Policy Workspace to Create and Modify Operational Policies

You can use the workflow in the policy workspace to create local policies quickly, and update the settings in existing policies. Select a base policy to use as the source for your local policy settings, and modify the thresholds and settings used for analysis and collection of data from objects or object groups in your environment. A policy that has no local settings defined inherits the settings from its base policy to apply to the associated objects or object groups.

Verify that objects or object groups exist for VMware Aria OperationsVMware Cloud Foundation Operations to analyze and collect data, and if they do not exist, create them. See Managing Custom Object Groups in .

1. From the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile.
2. Click **Add** to add a policy or you can select a policy and click **Edit Policy** to edit an existing policy.

You can add and edit policies and remove certain policies. You can use the Base Settings policy or the Default Policy as the root policy for the settings in other policies that you create. You can set any policy to be the default policy.

3. In the Create Policies workspace, assign a name to the policy, and enter the description.
   Give the policy a meaningful name and description so that all users know the purpose of the policy.
4. From the **Inherit From** drop-down, select one or more policies to use as a baseline to define the settings for your new local policy.
   You can use any of the policies provided with VMware Aria OperationsVMware Cloud Foundation Operations as a baseline source for your new policy settings.
5. Click **Create Policy**.

   The Create Policies workspace provides the options to customize your policy.

6. Click **Metrics and Properties**. In this workspace, select the metric, property, or super metric attributes to include in your policy.

   VMware Aria OperationsVMware Cloud Foundation Operations collects data from the objects in your environment based on the metric, property, or super metric attributes that you include in the policy.

   a) Click **Save** and return to the create policies workspace.
7. Click **Alerts and Symptoms**. In this workspace, select the alert definitions and symptom definitions, and activate or deactivate them as required for your policy.

   VMware Aria OperationsVMware Cloud Foundation Operations identifies problems on objects in your environment and triggers alerts when conditions occur that qualify as problems.

   a) Click **Save** and return to the create policies workspace.
8. Click **Capacity**. In this workspace, select and override the situational settings such as committed projects to calculate capacity, time remaining, and other detailed settings.
   a) Click **Save** and return to the create policies workspace.
9. Click **Compliance**. In this workspace, set the compliance threshold required for your policy.
   a) Click **Save** and return to the create policies workspace.
10. Click **Workload Automation**. In this workspace, select the optimization settings required for your policy.

    Click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option, your policy inherits the parent policy settings.

    a) Click **Save** and return to the create policies workspace.
11. Click **Groups and Objects**. In this workspace, select one or more groups and objects to which the policy applies.

    VMware Aria OperationsVMware Cloud Foundation Operations monitors the objects according to the settings in the policy that is applied to the object or the object group, triggers alerts when thresholds are violated, and reports the results in the dashboards, views, and reports. If you do not assign a policy to one or more objects or object groups, VMware Aria OperationsVMware Cloud Foundation Operations does not assign the settings in that policy to any objects, and the policy is not active. For an object or an object group that dos not have a policy assigned, VMware Aria OperationsVMware Cloud Foundation Operations associates the object group with the Default Policy.

    Filter the object types, and modify the settings for those object types so that VMware Aria OperationsVMware Cloud Foundation Operations collects and displays the data that you expect in the dashboards and views.

    a) Click **Save** and return to the create policies workspace.

After VMware Aria OperationsVMware Cloud Foundation Operations analyzes and collects data from the objects in your environment, review the data in the dashboards and views. If the data is not what you expected, edit your local policy to customize and override the settings until the dashboards display the data that you need.

### Policy Workspace in VMware Aria OperationsVMware Cloud Foundation Operations

The policy workspace allows you to quickly create and modify policies. To create a policy, you can inherit the settings from an existing policy, and you can modify the settings in existing policies if you have adequate permissions. After you create a policy, or edit an existing policy, you can apply the policy to one or more objects or object groups.

#### How the Policy Workspace Works

Every policy includes a set of packages, and uses the defined problems, symptoms, metrics, and properties in those packages to apply to specific objects or object groups in your environment. You can view the details for the settings inherited from the base policy, and display specific settings for certain object types. You can override the settings of other policies, and include additional policy settings to apply to the object types.

Use the **Add** and **Edit** options to create policies and edit existing policies.

#### Where You Create and Modify a Policy

To create and modify policies, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit the policy. The policy workspace is where you select the base policies, and customize and override the settings for analysis, metrics, properties, alert definitions, and symptom definitions. In this workspace, you can apply the policy to objects or object groups.

To remove a policy from the list, select the policy, click the horizontal ellipse, and select **Delete**.

#### Policy Workspace Options

The policy workspace includes a step-by-step workflow to create and edit a policy, and apply the policy to custom object groups.

### Getting Started Details

When you create a policy, you must give the policy a meaningful name and description so that users know the purpose of the policy.

#### Where You Assign the Policy Name and Description

To add a name and description to a policy, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. The name and description appear in the Create or Edit policy workspace.

**Table 31: Name and Description Options in the Create or Edit Policy Workspace**

| Option | Description |
|---|---|
| Name | Name of the policy as it appears in the Create or Edit Policy screens, and in areas where the policy applies to objects, such as Custom Groups. |
| Description | Meaningful description of the policy. For example, use the description to indicate which policy is inherited, and any specific information that users must understand the relationship of the policy to one or more groups of objects. |
| Inherit From | The base policy that is used as a starting point. All settings from the base policy will be inherited as default settings in your new policy. You can override these settings to customize the new policy. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Select a base policy to inherit the policy settings as a starting point for your new policy. |

**Select the Inherited Policy Details**

You can use any of the policies provided with VMware Aria OperationsVMware Cloud Foundation Operations as a baseline source for your policy settings when you create a policy.

In the policy content area, you can perform the following actions:
- View the packages and elements for the inherited policy and additional policies that you selected to override the settings.
- Compare the differences in settings highlighted between these policies.
- Display object types.

To create a policy, select a base policy to inherit your new custom policy inherits settings. To override some of the settings in the base policy according to the requirements for the service level agreement for your environment, you can select and apply a separate policy for a management pack solution. The override policy includes specific settings defined for the types of objects to override, either manually or that an adapter provides when it is integrated with VMware Aria OperationsVMware Cloud Foundation Operations. The settings in the override policy overwrite the settings in the base policy that you selected.

When you select and apply a policy to use to overwrite the settings that your policy inherits from the base policy, the policy that you select appears in the policy settings cards.

Click each card to display the inherited policy configuration, and your policy, and displays a preview of the selected policy settings. When you select one of the policy cards, you can view the number of activated and deactivated alert definitions, symptom definitions, metrics and properties, and the number of activated and deactivated changes.

When you select the Groups and Objects card, you select the objects to view so that you can see which policy elements apply to the object type. For example, when you select the StorageArray object type, the workspace displays the local packages for the policy and the object group types with the number of policy elements in each group.

You can preview the policy settings for all object types, only the object types that have settings changed locally, or settings for new object types that you add to the list, such as Storage Array storage devices.

**Where You Select and Override Base Policies Settings**

To select a base policy to use as a starting point for your own policy, and to select a policy to override one or more settings that your policy inherits from the base policy, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy. In the Create policies workspace, add a name and description for the policy and from the **Inherit From** drop-down, select the base policy. The policy configuration, objects, and preview appear in cards below this drop-down.

**Capacity Details**

You can filter the object types, and modify the settings for those object types so that VMware Aria OperationsVMware Cloud Foundation Operations applies these settings. The data that you expect then appears in the dashboards and views.

## How the Capacity Workspace Works

When you turn on and configure the Capacity settings for a policy, you can override the settings for the policy elements that VMware Aria OperationsVMware Cloud Foundation Operations uses to trigger alerts and display data. These types of settings include symptom thresholds based on alerts, situational settings such as committed projects to calculate capacity and time remaining, and other detailed settings.

Policies focus on objects and object groups. When you configure policy settings for your local policy, you must consider the object type and the results that you expect to see in the dashboards and views. If you do not change these settings, your local policy retains the settings that your policy inherited from the base policy that you selected.

## Where You Set the Policy Capacity Settings

To set the capacity settings for your policy, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy.

In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The capacity settings for host systems, virtual machines, and other object types that you select appears in the workspace.

You can also edit the capacity settings while working on the objects under the Environment Tab. In the **Capacity** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Capacity Setting**.

**Table 32: Capacity Settings in the Create or Edit Policy Workspace**

| Option | Description |
|---|---|
| Risk Level Configurations | You can set the risk level for the time that is remaining when the forecasted total need of a metric reaches usable capacity. Click the lock icon to override the settings and change the thresholds for your policy.<br><br>The following are the risk level settings. Use the slider below the graphical display to change the risk level. You can move the slider between Aggressive and Conservative.<br><br>• Conservative. Use this option for production and mission-critical workloads.<br>• Aggressive. Use this option for non-critical workloads.<br>• Peak focused. Selecting peak focused tells the capacity engine to create projections using the peaks that have been identified in the historical demand. Use this option to include the upper range of the data. The projection will be based on the high utilization points. Select the Peak focused checkbox for VMs with utilization spikes. |
| Business Hours Schedule | Configure business hours as per your time zone, for calculation of capacity analysis and projections. VMware Aria OperationsVMware Cloud Foundation Operations considers the business hours for all objects using the current policy.<br><br>During non-business hours, VMs could be running other data center activities such as OS upgrades, virus scans, etc after working hours, and hence may not appear to be idle. When you mark business hours schedules, VMware Aria OperationsVMware Cloud Foundation Operations can analyze after hours metrics for inventory, compliance, troubleshooting and other purposes. The reclamation and right sizing analysis and recommendations are based on the business hours and ignore spikes after business hours.<br><br>Since the business hours schedule are based on policies, different objects can have different business hours. The capacity charts will be based on business hours. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>You can set business hours schedule for VMs and clusters only.<br><br>**NOTE**<br>After you specify business hours, the capacity forecast for the object will be based on the business hours and not 24 hours.<br><br>Click the lock icon on the left of each element to override the settings and change the thresholds for your policy. |
| Filters | Select the object type by which you want to filter. You can filter by Object Types, Local Changes, and Unsaved Changes. |
| Capacity Settings | Select an object to view the policy elements and settings for the object type so that you can have VMware Aria OperationsVMware Cloud Foundation Operations analyze the object type.<br><br>You can view and modify the settings for the following policy elements:<br>• Allocation Model<br>• Custom Profile<br>• Capacity Buffer<br><br>Click the lock icon on the left of each element to override the settings and change the thresholds for your policy. |
| Criticality Thresholds and Metrics | There are two tabs in this settings.<br><br>Click the lock icon on the left of each element to override the settings and change the thresholds for your policy.<br><br>**Criticality Thresholds Tab**<br>You can view and modify the threshold settings for the following policy elements:<br>• Time Remaining<br>• Capacity Remaining<br>• Workload<br>**Custom Metrics Tab**<br>In the custom metrics tab, you can configure VMware Aria Operations to use custom metrics in all the capacity calculations. The metrics that you configure in this tab replaces the default metrics that the VMware Aria Operations capacity engine uses. When defining the custom metrics, you can select the metrics shipped with VMware Aria Operations, or select super metrics. Only metrics which have the same unit as the internal metric used by the capacity engine, or a metric which has no unit, are displayed.<br>**NOTE**<br>Enabling custom metrics in the capacity calculations is an advanced configuration. Custom metrics alter the way VMware Aria Operations calculates capacity across your environment. Use this setting only when needed. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | You can view and modify the custom metrics settings for all non-allocation capacity models. For example, for a data center, you can set custom metrics settings for Total Capacity and Utilization for Memory, CPU and Demand.<br>When you click the **Edit** icon beside the total Capacity and Utililization settings, a list of available metrics opens in the right pane. Double click a default metric or super metric from the list to select it. Click **RESET TO DEFAULT** to revert your changes. Changes that you make take effect after the next collection cycle. |

Click **Save** to save the changes.

The local changes made will appear under **Policies › Default Policy › Capacity** section. You can also view the preview of changes in the Capacity card.

### Policy Workload Element

Workload is a measurement of the demand for resources on an object. You can turn on and configure the settings for the Workload element for the object types in your policy.

How the Workload Element Works

The Workload element determines how VMware Aria OperationsVMware Cloud Foundation Operations reports on the resources that the selected object group uses. The resources available to the object group depend on the amount of configured and usable resources.

- A specific amount of physical memory is a configured resource for a host system, and a specific number of CPUs is a configured resource for a virtual machine.
- The usable resource for an object or an object group is a subset of, or equal to, the configured amount.
- The configured and usable amount of a resource can vary depending on the type of resource and the amount of virtualization overhead required, such as the memory that an ESX host machine requires to run the host system. When accounting for overhead, the resources required for overhead are not considered to be usable, because of the reservations required for virtual machines or for the high availability buffer.

Where You Override the Policy Workload Element

To view and override the policy workload capacity setting, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The workload settings for the object type that you have selected appear in the workspace.

View the Workload policy element, and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

**Table 33: Policy Workload Element Settings in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Workload | Allows you to set the number of collection cycles it takes to trigger or clear an alert. |

### Policy Time Remaining Element

The Time remaining element is a measure of the amount of time left before your objects run out of capacity.

How the Time Remaining Element Works

The Time Remaining element determines how VMware Aria OperationsVMware Cloud Foundation Operations reports on the available time until capacity runs out for a specific object type group.

- The time remaining indicates the amount of time that remains before the object group consumes the capacity available. VMware Aria OperationsVMware Cloud Foundation Operations calculates the time remaining as the number of days remaining until all the capacity is consumed.
- To keep the Time Remaining more than the critical threshold setting or to keep it green, your objects must have more days of capacity available.

Where You Override the Policy Time Remaining Element

To view and override the policy Time Remaining capacity setting, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The time remaining settings for the object type that you have selected appear in the workspace.

View the Time Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

**Table 34: Policy Time Remaining Element Settings in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Time Remaining | Allows you to set the number of days until capacity is projected to run out based on your current consumption trend. |

### Policy Capacity Remaining Element

Capacity is a measurement of the amount of memory, CPU, and disk space for an object. You can turn on and configure the settings for the Capacity Remaining element for the object types in your policy.

How the Capacity Remaining Element Works

The Capacity Remaining element determines how reports on the available capacity until resources run out for a specific object type group.

- The capacity remaining indicates the capability of your environment to accommodate workload.
- Usable capacity is a measurement of the percentage of capacity available, minus the capacity affected when you use high availability.

Where You Override the Policy Capacity Remaining Element

To view and override the policy Capacity Remaining analysis setting, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The capacity remaining settings for the object type that you have selected appears in the workspace.

View the Capacity Remaining policy element and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

**Table 35: Policy Capacity Remaining Element Settings in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Capacity Remaining | Allows you to set the percentage at which the capacity remaining alerts must be triggered. |

### Policy Allocation Model Element

Allocation model defines how much CPU, memory, or disk space is allocated to objects in a datastore, cluster or datastore cluster. In the policy, you can turn on the Allocation Model element and configure the resource allocation for the objects.

How the Allocation Model Element Works

The Allocation Model element determines how calculates capacity when you allocate a specific amount of CPU, memory, and disk space resource to datastores, clusters or datastore clusters. You can specify the allocation ratio for either one, or all of the resource containers of the cluster. Unlike the demand model, the allocation model is used for capacity calculations only when you turn it on in the policy.

The allocation model element also affects the reclaimable resources for memory and storage in Reclaim page. When you turn on the Allocation Model element in the policy, the tabular representation of the VMs and snapshots in the selected data center from which resources can be reclaimed displays reclaimable memory and disk space based on the overcommit values.

Where You Override the Allocation Model Element

To view and override the policy workload analysis setting, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile.

Click **Add** to add a policy or select the required policy, and then in the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card.

The allocation model settings for the object type that you selected appear in the workspace.

Click the unlock icon next to Allocation Model to set the overcommit ratios.

| Option | Description |
|---|---|
| Set overcommit ratio, to enable Allocation Model | Allows you to set the overcommit ratio for CPU, memory, or disk space. Select the check box next to the resource container you want to edit and change the overcommit ratio value. |

### Policy Custom Profile Element

The custom profile element lets you apply a custom profile which shows how many more of a specified object can fit in your environment depending on the available capacity and object configuration.

Where You Define the Custom Profiles

To define a custom profile, from the left menu click **Operations › Configurations**, and then click the **Custom Profiles** tile under Miscellaneous. Click **Add** to define a new custom profile.

Where You Select the Custom Profile Element

To view and override the policy Custom Profile analysis setting, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The custom profile element for the object types such as datastores, clusters and datastore clusters that you selected appear in the workspace. Click the lock icon to unlock the section and make changes.

### Policy Capacity Buffer Element

The capacity buffer element lets you add buffer for capacity and cost calculation. For Key definition for "vCenter_server" not found in the DITA map. objects, you can add buffer to CPU, Memory, and Disk Space for the Demand and Allocation models. You can add capacity buffer to datastores, clusters and datastore clusters. The values that you define here affect the cluster cost calculation. The time remaining, capacity remaining, and recommended values are calculated based on the buffer. For WLP, capacity buffer is first considered and then the headroom that you have defined is considered.

Where You Define the Capacity Buffer

To view and override the policy Capacity Buffer analysis setting, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the <policy name> [Edit] workspace, click the **Capacity** card. The Capacity Buffer for the object type that you selected appears in the workspace. Click the lock icon to unlock the section and make changes.

**How the Capacity Buffer Element Works**

The Capacity Buffer element determines how much extra headroom you have and ensures that you have extra space for growth inside the cluster when required. The value of the usable capacity reduces by the buffer amount that you specify here. The default buffer value is zero. If you are upgrading from a previous version of VMware Aria OperationsVMware Cloud Foundation Operations, the buffer values are carried forward to the new version.

The capacity buffer value that you specify for the Allocation model is considered only if you have activated allocation model in the policy.

Starting from version 8.6, capacity buffer is depreciated from cluster compute resources. The overcommit ratio setting (from the allocation model) and buffer settings, if set for the datastore object, takes precedence for the disk space related to datastore cluster and cluster objects. If these settings are not set, then, from a cost calculation perspective, the settings of datastore cluster and cluster (if the settings are missing for the datastore cluster as well), are used. The allocation and buffer settings made on the cluster does not impact the underlying datastores (as they do not inherit these settings), and the same works vice-versa, settings made for datastores are not propagated to the cluster.

The following tables display the capacity buffer that you can define based on the vCenter Adapter object types:

| Object Type | Valid Models for Capacity Buffer |
|---|---|
| CPU | Demand<br>Allocation |
| Memory | Demand<br>Allocation |
| Disk Space | Demand<br>Allocation |

### Maintenance Schedule Details

You can set a time to perform maintenance tasks for each policy.

### Where You Override the Policy Maintenance Schedule Element

To view and override the policy Maintenance Schedule analysis setting, from the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Maintenance Schedule**.
If you do not configure the policy element, your policy inherits the settings from the selected base policy.

**Table 36: Policy Maintenance Schedule Element Settings in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Select Object Type | Select the object type by which you want to filter. |
| Filters | You can filter by Local Changes and Unsaved Changes. Select **Yes** or **No** from the drop-down and click **Apply** to apply the filters. |
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Maintenance Schedule | Sets a time to perform maintenance tasks. During maintenance, VMware Aria OperationsVMware Cloud Foundation Operations does not calculate analytics. |

## Compliance Details

Compliance is a measurement that ensures that the objects in your environment meet industrial, governmental, regulatory, or internal standards. You can unlock and configure the settings for the compliance for the object types in your policy.

### Where You Override the Policy Compliance

To view and override the policy compliance setting, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Compliance**
View the compliance thresholds and configure the settings for your policy.

If you do not configure the policy element, your policy inherits the settings from the selected base policy.

**Table 37: Compliance Settings in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Lock icon | Allows you to override the policy element settings so that you can customize the policy to monitor the objects in your environment. |
| Compliance | Allows you to set the compliance score threshold based on the number of violations against those standards. |

## Workload Automation Details

You can set the workload automation options for your policy, so that VMware Aria OperationsVMware Cloud Foundation Operations can optimize the workload in your environment as per your definition.

### How the Workload Automation Workspace Works

You click the lock icon to unlock and configure the workload automation options specific for your policy. When you click the lock icon to lock the option, your policy inherits the parent policy settings.

**Where You Set the Policy Workload Automation**

Access this screen through the Policies pages:

1. Click **Operations › Configurations**, and then click the **Policy Definition** tile.
2. Select a policy that you want to modify. Ideally, this should be an active policy. Or, click the **ADD** button to add a new policy.
3. Select the **Workload Automation** card to review the changes, or click **EDIT POLICY** to make changes.

**Table 38: Workload Automation in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Workload Optimization | Select a goal for workload optimization.<br>Select **Balance** when workload performance is your first goal. This approach proactively moves workloads so that the resource utilization is balanced, leading to maximum headroom for all resources.<br><br>Select **Moderate** when you want to minimize the workload contention.<br><br>Select **Consolidate** to proactively minimize the number of clusters used by workloads. You might be able to repurpose resources that are freed up. This approach is good for cost optimization, while making sure that performance goals are met. This approach might reduce licensing and power costs. |
| Cluster Headroom | Headroom establishes a required capacity buffer, for example, 20 percent. It provides you with an extra level of control and ensures that you have extra space for growth inside the cluster when required. Defining a large headroom setting limits the systems opportunities for optimization.<br><br>**NOTE**<br>vSphere HA overhead is already included in useable capacity and this setting does not impact the HA overhead. |
| Change Datastore | Click the lock icon to select one of the following options:<br>• Do not allow Storage vMotion.<br>• Allow Storage vMotion. This is selected by default.<br>Using this option, you can select what type of virtual machines VMware Aria OperationsVMware Cloud Foundation Operations moves first to address workload. |
| Target Network Policy Setting for WLP | Click the lock icon to select the following option:<br>• Generate a Target Network mapping<br>When you select this checkbox, the Workload Placement algorithm in VMware Cloud Foundation Operations will automatically choose compatible target network, while making the decision to move the VM for the optimization. For choosing compatible network WLP engine will consider the segment path and logical switch UUID of the Distributed Port Group.<br><br>Workload Optimization Across networks is supported when the optimization candidate clusters are assigned with different port groups (configured with NSX). These port groups configured via NSX have same segmentID and Logical Switch UUID. To enable this ability, check the respective setting in VMware Cloud Foundation Operations Workload Automation policy settings. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>Segment ID and logical switch UUID properties are published on the VC port groups by NSX. So Worload Placement cannot provide a target network if it is not a NSX configuration and those properties are missing.<br><br>This setting is not selected by default. |

**Configuring vCenter Pricing Details**

You can add and assign new pricing cards to vCenter and Clusters in VMware Aria OperationsVMware Cloud Foundation Operations. The pricing card can be cost-based or rate-based, you can customize the cost-based pricing card and rate-based pricing card as per your requirement. After configuring the pricing card, you can assign it to one more vCenter or Clusters based on your pricing strategy.

If you want to copy the vCenter pricing settings from the policy currently being edited to another policy, click **Copy local changes to other policy** and select the policy to which you want to copy the settings. The copied pricing configuration will override any existing local pricing configuration, in the target policy.

1. From the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile.
2. Select the required policy or click **Add** to add a new policy.
3. In the right pane, click **Edit Policy**.
4. In the <policy name> [Edit] workspace, click the **VC Pricing** card.
5. Click the Lock icon to override parent policy settings.
6. Select if you want to activate or deactivate the pricing engine.
7. Configure **Basic Charges**: Click the Lock icon to edit the parent policy settings. Pricing can be performed either on a cost basis or independent of it by specifying rate cards. The factor entered here is multiplied by the cost calculated as a derivative of cost drivers.

    1. Based on Cost/Based on Rate: Select if you want to pricing card to be cost-based or rate-based.
       The following options appear if you select the **Based on Cost** option:
       - CPU Cost: Enter a valid CPU cost factor.
       - Memory Cost: Enter a valid memory cost factor.
       - Storage Cost: Enter a valid storage cost factor.
       - Additional Cost: Enter a valid additional cost factor.

       The following options appear if you select the **Based on Rate** option:
       - CPU Rate: Enter the CPU Rate per vCPU, the charging period, and how to charge for the resources.
       - Memory Rate: Enter the memory rate per GB, the charging period, and how to charge for the resources.
       - Storage Rate: Enter the storage rate per GB, the charging period, and how to charge for the resources.

8. Configure **Guest OS Rate**: Click **Guest OS Rate** in the left pane and then click the Lock icon to edit the parent policy settings. These are additional charges that have to be included based on the operating system running on the virtual machine. The name of the operating system should match exactly as discovered by VMware Tools.

    1. Click **Create Guest OS Rate** and enter the following details:
       – Guest OS Name: Enter a guest OS name.

- Charge Period: The Charge Period indicates the frequency of charging.
- Base Rate: Enter a base rate.

The guest OS rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

9. **Tags**: Click **Tags** in the left pane and then click the Lock icon to edit the parent policy settings. Tag-based charges can be used to charge for value-added services such as antivirus database disaster recovery and other applications. These applications are to be represented as vCenter tags on the VMs for these charges to work.

    1. Recurring Charges: Recurring charges represent repeating charges such as monthly license fees for antivirus software. Click **Add Recurring Tag** and enter the following details:
        - Tag Category: Enter a tag key.
        - Tag Value: Enter a tag value.
        - Base Rate: Enter a base rate.
        - Charge Period: The Charge Period indicates the frequency of charging.
        - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.

        The tags that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
    2. One Time Tag: Tag-based one-time charges can be used to represent incidental charges such as charges for addressing a support ticket or charges for applying an operating systems patch. Click **Add One Time Tag** and enter the following details:
        - Tag Category: Enter a tag key.
        - Tag Value: Enter a tag value.
        - Base Rate: Enter a base rate.

        The tags that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
    3. Rate Factor Tag: Rate factors are multiplication factors applied to already calculated charges. For example, to add a 50% premium on storage, set a rate factor of 1.5 to storage charge. Click **Rate Factor Tag** and enter the following details:
        - Tag Category: Enter a tag key.
        - Tag Value: Enter a tag value.
        - Charge Applies To: Select what the charge applies to.
        - Rate Factor: Enter a valid number. For example, if you want to increase the price of CPU which has a tag 'Tag1-Value1' by 20% then select `CPU Charge` from the **Charge Applies To** drop-down list and enter `1.2` in **Rate Factor**.

        The tags that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

10. Configure **Overall Charges**: Click **Overall Charges** in the left pane and then click the Lock icon to edit the parent policy settings. Overall Charges are flat charges that are applied to VMs that match this policy.

    1. VM Setup Charges: Enter a valid setup fee. This is to charge for the setup of the VMs.
    2. Recurring Charges: Enter a valid number.
    3. Charge Period: The Charge Period indicates the frequency of charging.

You can assign policies to the required Organization/Organization VDC under **Operations › Configurations**, and then click the **Policy Assignment** tile. For details, see Assigning Policies.

### Metrics and Properties Details

You can select the attribute type to include in your policy so that VMware Aria OperationsVMware Cloud Foundation Operations can collect data from the objects in your environment. Attribute types include metrics, properties, and super metrics. You activate or deactivate each metric, and determine whether to inherit the metrics from base policies that you selected in the workspace.

#### How the Collect Metrics and Properties Workspace Works

When you create or customize a policy, you can override the base policy settings to have VMware Aria OperationsVMware Cloud Foundation Operations collect the data that you intend to use to generate alerts, and report the results in the dashboards.

To define the metric and super metric symptoms, metric event symptoms, and property symptoms, from the left menu click **Operations › Configurations**, and then click the **Symptom Definitions** tile.

#### Where You Override the Policy Attributes

To override the attributes and properties settings for your policy, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policy workspace, click **Metrics and Properties**. The attributes and properties settings for the selected object types appear in the workspace.

You can also edit the metrics and properties while working on the objects under the Environment Tab. In the **Metrics** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Metrics Collection**.

**Table 39: Metrics and Properties Options**

| Option | Description |
|---|---|
| Actions | Select one or more attributes and select activate, deactivate, or inherit to change the state and KPI for this policy. |
| Filter options | Deselect the options in the **Attribute Type**, **State**, **KPI**, and **DT** drop-down menus, to narrow the list of attributes.<br>• ✔ Activated. Indicates that an attribute will be calculated.<br>• ✔ Activated (Force). Indicates state change due to a dependency.<br>• ⊘ Deactivated. Indicates that an attribute will not be calculated.<br>• ✔ Inherited. Indicates that the state of this attribute is inherited from the base policy and will be calculated.<br>• ⊘ Inherited. Indicates that the state of this attribute is inherited from the base policy and will not be calculated.<br><br>The KPI determines whether the metric, property, or super metric attribute is considered to be a key performance indicator (KPI) when VMware Aria OperationsVMware Cloud Foundation Operations reports the collected data in the dashboards. Filter the KPI states to display attributes with KPI activated, deactivated, or inherited for the policy. |
| Object Type | Filters the attributes list by object type. |
| Page Size | The number of attributes to list per page. |
| Attributes data grid | Display the attributes for a specific object type.<br>• Name. Identifies the name of the metric or property for the selected object type. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Type. Distinguishes the type of attribute to be either a metric, property, or super metric.<br>• Adapter Type. Identifies the adapter used based on the object type selected, such as Storage Devices.<br>• Object Type. Identifies the type of object in your environment, such as StorageArray.<br>• State. Indicates whether the metric, property, or super metric is inherited from the base policy.<br>• KPI. Indicates whether the key performance indicator is inherited from the base policy. If a violation against a KPI occurs, VMware Aria OperationsVMware Cloud Foundation Operations generates an alert.<br>• DT. Indicates whether the dynamic threshold (DT) is inherited from the base policy. |

**Alert and Symptom Details**

You can activate or deactivate alert and symptom definitions to have VMware Aria OperationsVMware Cloud Foundation Operations identify problems on objects in your environment and trigger alerts when conditions occur that qualify as problems. You can automate alerts.

**How the Alert and Symptom Definitions Workspace Works**

VMware Aria OperationsVMware Cloud Foundation Operations collects data for objects and compares the collected data to the alert definitions and symptom definitions defined for that object type. Alert definitions include associated symptom definitions, which identify conditions on attributes, properties, metrics, and events.

You can configure your local policy to inherit alert definitions from the base policies that you select, or you can override the alert definitions and symptom definitions for your local policy.

Before you add or override the alert definitions and symptom definitions for a policy, familiarize yourself on the available alerts and symptoms.

- To view the available alert definitions, from the left menu, click **Operations › Configurations**, and then click the **Alert Definitions** tile.
- To view the available symptom definitions, from the left menu, click **Operations › Configurations**, and then click the **Symptom Definitions** tile. Symptom definitions are available for metrics, properties, messages, faults, smart early warnings, and external events.

A summary of the number of problem and symptoms that are activated and deactivated, and the difference in changes of the problem and symptoms as compared to the base policy, appear in the Analysis Settings pane of the policies workspace.

**Where You Override the Alert Definitions and Symptom Definitions**

To override the alert definitions and symptom definitions for your policy, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The definitions appear in the workspace.

You can also edit the alert settings while working on the objects under the Environment Tab. In the **Alerts** tab under **Environment**, click the **Foundation Policy** drop-down and select **Edit Alerts State**.

**Policy Alert Definitions and Symptom Definitions**

You can override the alert definitions and symptom definitions for each policy.

## Policy Alert Definitions

Each policy includes alert definitions. Each alert uses a combination of symptoms and recommendations to identify a condition that classifies as a problem, such as failures or high stress. You can activate or deactivate the alert definitions in your policy, and you can set actions to be automated when an alert triggers.

How the Policy Alert Definitions Work

VMware Aria OperationsVMware Cloud Foundation Operations uses problems to trigger alerts. A problem manifests when a set of symptoms exists for an object, and requires you to take action on the problem. Alerts indicate problems in your environment. VMware Aria OperationsVMware Cloud Foundation Operations generates alerts when the collected data for an object is compared to alert definitions for that object type and the defined symptoms are true. When an alert occurs, VMware Aria OperationsVMware Cloud Foundation Operations presents the triggering symptoms for you to take action.

Some of the alert definitions include predefined symptoms. When you include symptoms in an alert definition, and activate the alert, an alert is generated when the symptoms are true.

The Alert Definitions pane displays the name of the alert, the number of symptoms defined, the adapter, object types such as host or cluster, and whether the alert is activated as indicated by **Local**, deactivated as indicated by **not Local**, or inherited. Alerts are inherited with a green checkmark by default, which means that they are activated.

You can automate an alert definition in a policy when the highest priority recommendation for the alert has an associated action.

To view a specific set of alerts, you can select the badge type, criticality type, and the state of the alert to filter the view. For example, you can set the policy to send fault alerts for virtual machines.

Where You Modify the Policy Alert Definitions

To modify the alerts associated with policies, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

**Table 40: Alert Definitions in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Object Type | Filters the alert definitions list by object type. |
| Filters | Limits the list based on the text you type.<br><br>You can also filter by:<br>• Name<br>• Criticality<br>• Impact<br>• State<br>• Automate<br>• Local Changes<br>• Unsaved Changes<br><br>Impact indicates the health, risk, and efficiency badges to which the alerts apply.<br><br>Criticality indicates the information, critical, immediate, warning, or automatic criticality types to which the alert definition applies. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Automate indicates the actions that are activated for automation when an alert triggers, or actions that are deactivated or inherited. Actions that are activated for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to **Local** with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark. |
| Actions | Select one or more alert definitions and select activate, deactivate, or inherit to change the state for this policy. |
| Page Size | The number of alert definitions to list per page. |
| Alert Definitions data grid | Displays information about the alert definitions for the object types. The full name for Alert definition and the criticality icon appear in a tooltip when you hover the mouse over the Alert Definition name.<br>• Alert Definition. Meaningful name for the alert definition.<br>• State. Alert definition state, either activated, deactivated, or inherited from the base policy.<br>• Automate. When the action is set to Local, the action is activated for automation when an alert triggers. Actions that are activated for automation might appear as inherited with a green checkmark, because policies can inherit settings from each other. For example, if the Automate setting in the base policy is set to **Local** with a green checkmark, other policies that inherit this setting will display the setting as inherited with a green checkmark.<br>• Symptom. Number of symptoms defined for the alert.<br>• Criticality. Indicates the criticality of the alert.<br>• Actionable Recommendations. Only recommendations with actions in the first priority, as they are the only ones you can automate.<br>• Adapter. Data source type for which the alert is defined.<br>• Object Type. Type of object to which the alert applies. |

If you do not configure the package, the policy inherits the settings from the selected base policy.

## Activating the Deactivated Alerts

Several out-of-the-box alerts have been deactivated to enhance your alert experience and reduce alert noise in your environment. The alerts that are triggered for these deactivated alerts are auto-cancelled and as a result, you may experience a dip in the number of alerts triggered. However, you can still activate these alerts in specific policies.

The reason for deactivating these alerts is that there could be an overwhelming number of alerts when alerts are turned on for all objects, making it difficult to identify the ones that need immediate attention. It is recommended to exercise caution while activating the deactivated alerts for applicable policies.

Read the KB article, KB 91410 to know the list of deactivated alerts.

Perform the following steps to activate the deactivated alerts:

1. From the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile.
2. Select the required policy and in the right pane, click **Edit Policy**, and then select the **Alerts and Symptoms** tile.
3. Go to Filters and enter the name of the deactivated alert, and click **Apply**. You can refer to the KB article, KB 91410 for the list of deactivated alerts.
4. Select **Activated** from the **State** drop-down list or click **Actions › State › Activated**.
5. Click **Save**.

**NOTE**
You can also activate all the deactivated alerts at once. To do this, filter the alerts by **Deactivated** State, click on the **Select All** option, and from the **Actions** drop-down list, click **State › Activated**.

**NOTE**
You can also activate the deactivated alerts by creating a separate policy, adding custom groups in that policy, and activating the alert definitions in the required policy. By doing this, the deactivated alerts are activated in the user-defined policy and will apply to the objects in the custom group. For more details on custom groups, see Managing Custom Object Groups in VMware Aria Operations.

The deactivated alerts are now active.

### Policy Symptom Definitions

Each policy includes a package of symptom definitions. Each symptom represents a distinct test condition on a property, metric, or event. You can activate or deactivate the symptom definitions in your policy.

How the Policy Symptom Definitions Work

VMware Aria OperationsVMware Cloud Foundation Operations uses symptoms that are activated to generate alerts. When the symptoms used in an alert definition are true, and the alert is activated, an alert is generated.

When a symptom exists for an object, the problem exists and requires that you take action to solve it. When an alert occurs, VMware Aria OperationsVMware Cloud Foundation Operations presents the triggering symptoms, so that you can evaluate the object in your environment, and with recommendations for how to resolve the alert.

To assess objects for symptoms, you can include symptoms packages in your policy for metrics and super metrics, properties, message events, and faults. You can activate or deactivate the symptoms to determine the criteria that the policy uses to assess and evaluate the data collected from the objects to which the policy applies. You can also override the threshold, criticality, wait cycles, and cancel cycles.

The Symptoms pane displays the name of the symptom, the associated management pack adapter, object type, metric or property type, a definition of the trigger such as for CPU usage, the state of the symptom, and the trigger condition. To view a specific set of symptoms in the package, you can select the adapter type, object type, metric or property type, and the state of the symptom.

When a symptom is required by an alert, the state of the symptom is activated, but is dimmed so that you cannot modify it. The state of a required symptom includes an information icon that you can hover over to identify the alert that required this symptom.

Where You Modify the Policy Symptom Definitions

To modify the policy package of symptoms, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Alerts and Symptoms**. The alert definitions and symptom definitions for the selected object types appear in the workspace.

**Table 41: Symptom Definitions in the Create or Edit Policies Workspace**

| Option | Description |
|---|---|
| Object Type | Select an object type to view the symptom definitions list by the selected object type. |
| Filters | Limits the list based on the text you type.<br><br>You can also filter by:<br>• Name<br>• Criticality<br>• Type<br>• State<br>• Local Changes<br>• Unsaved Changes |
| Actions | Select one or more symptom definitions and select activate, deactivate, or inherit to change the state for this policy. |
| Page Size | The number of symptom definitions to list per page. |
| Symptom Definitions data grid | Displays information about the symptom definitions for the object types. The full name for Symptom Definition appears in a tooltip when you hover the mouse over the Symptom Definition name.<br>• Symptom Definition. Symptom definition name as defined in the list of symptom definitions in the Content area. Click this name to view the details of the symptom.<br>• State. Symptom definition state, either activated, deactivated, or inherited from the base policy.<br>– ✔ Activated. Indicates that a symptom definition will be included.<br>– ✔ Activated (Force). Indicates state change due to a dependency.<br>– ⊘ Deactivated. Indicates that a symptom definition not be included.<br>– ✔ Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will be included.<br>– ⊘ Inherited. Indicates that the state of this symptom definition is inherited from the base policy and will not be included.<br>• Threshold. To change the threshold, you must set the State to **Activated**, set the condition to **Override**, and set the new threshold in the Override Symptom Definition Threshold dialog box.<br>• Type. Type of object to which the alert applies. Type determines whether symptom definitions that apply to HT and DT metrics, properties, events such as message, fault, and metric, and smart early warnings appear in the list.<br>• Criticality. Indicates the criticality.<br>• Adapter. Data source type for which the alert is defined.<br>• Object Type. Object type on which the symptom definition must be evaluated.<br>• Trigger. Static or dynamic threshold, based on the number of symptom definitions, the object type and metrics selected, the numeric value assigned to the symptom definition, the criticality of the symptom, and the number of wait and cancel cycles applied to the symptom definition.<br>• To activate the metric threshold evaluation based on real-time data collected every 20 seconds, click the **Near Real-Time Monitoring** check box.<br>• Condition. Activates action on the threshold. When set to Override, you can change the threshold. Otherwise set to default. |

If you do not configure the package, the policy inherits the settings from the selected base policy.

### Groups and Objects details

You can assign your local policy to one or more objects or groups of objects to have VMware Aria OperationsVMware Cloud Foundation Operations analyze those objects according to the settings in your policy. You can trigger alerts when the defined threshold levels are violated, and display the results in your dashboards, views, and reports.

#### How the Groups and Objects Workspace Works

When you create a policy, or modify the settings in an existing policy, you apply the policy to one or more objects or groups of objects. VMware Aria OperationsVMware Cloud Foundation Operations uses the settings in the policy to analyze and collect data from the associated objects, and displays the data in dashboards, views, and reports.

#### Where You Apply a Policy to Groups and Objects

To apply the policy to an object or groups of objects, from the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**.

#### Groups and Objects Options

To apply the policy to an object or groups of objects, select the check box for the groups or objects in the workspace.

You can then view the groups and objects associated with the policy. From the left menu click **Operations › Configurations**, and then click the **Policy Definition** tile. Click **Add** to add a policy or select the required policy. In the right pane, click **Edit Policy** to edit a policy. In the Create or Edit policies workspace, click **Groups and Objects**. Click the **Custom Groups** tab to apply the policy to one or more groups of objects. Click the **Objects** tabs to apply the policy to one or more objects.

For more information about how to create an object group, see the topic called **Custom Object Groups Workspace to Create a New Group**.

For more information about how to create a policy, see Policy Workspace in VMware Aria Operations.

## Assigning Policies

The **Policy Assignment** workspace displays all the policies available in your environment.

#### Where You Assign Policies

To create and modify policy assignments, from the left menu, click **Operations › Configurations**, and then click the **Policy Assignment** tile.

#### How the Policy Assignment Workspace Works

You can assign policies to your environment to activate controls, view, and manage your object assignment scope.

| Option | Description |
|---|---|
| Left Pane: Displays the list of policies available in your environment and provides different filters to understand policy distribution. | |
| Inactive/Active | You can filter policies by active or inactive status. A policy is considered to be active only when it is assigned to an object or a custom group. |
| Visualization Type | You can sort active policies based on different criteria. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Fixed: Displays active policies with the sections of the policies fixed and equal to each other.<br>• By Assignments: Displays the policies based on the number of objects assigned.<br>• By Affected Objects: Displays policies based on the affected object count.<br><br>**NOTE**<br>The sorting options are available only for the active policies. |
| Type Policy Name | Type the name of the policy in the search box to filter by policy name. |
| Policies | All the policies available in your environment are displayed in the left pane. Expand the policy card to view the priority, direct assignments, custom groups, affected objects, default assignment, and hierarchy of the policy.<br><br>**NOTE**<br>The default assignments are displayed only for the default policy.<br><br>When you click on a policy, the objects and custom groups associated with the selected policy are displayed under **Assigned Objects** in the right pane. |
| Right Pane: Displays details of the selected policy and allows you to drag and drop objects or custom groups to the selected policy. | |
| Selected Policy | Displays the name of the selected policy. |
| All Objects | Displays all objects and custom groups available in your environment in the **Inventory** and **Custom Groups** tabs respectively.<br>To add new objects to a policy:<br>1. In the right pane, click **All Objects › Inventory**.<br>2. Select the objects you want to assign to the policy. You can also search for the objects by typing the object name in the search box.<br>3. Drag the items from the **Inventory** tab and drop them into the policy card on the left pane.<br>4. In the **Assign Objects** window, select one of the following options:<br>a. Only this object: Select this if you want to apply changes only to the selected objects.<br>b. Include child object: Select this if you want to apply changes to the child objects. You can define the depth of change by entering a number in the **Depth** field. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>The maximum value for the Depth field is 10.<br>5. Click **Confirm**.<br><br>To add custom groups to a policy:<br>1. In the right pane, click **All Objects › Custom Groups**.<br>2. Select the custom groups you want to assign to the policy in the **Custom Groups** tab. You can also search for the custom groups by typing the custom group name in the search box.<br>3. Drag the items from the **Custom Groups** tab and drop them into the policy card on the left pane.<br>4. Click **Confirm**. |
| Assigned Objects | Displays the objects and custom groups assigned to the selected policy.<br>• Name: Displays the name of the object/custom group.<br>• Assignment Type: Displays the type of assignment.<br>• Depth: Displays the depth of the child objects that the policy affects. Click the Edit icon to change the depth.<br>• Action: Allows you to delete an object or custom group.<br><br>You can also drag objects/custom groups from **Assigned Objects** and drop them into the policy card of your choice. |

## Integrating Data Sources with VMware Aria OperationsVMware Cloud Foundation Operations

You can extend the monitoring capabilities of VMware Aria OperationsVMware Cloud Foundation Operations by installing and configuring integrations in VMware Aria OperationsVMware Cloud Foundation Operations to connect to, and analyze data from external data sources in your environment. Once connected, you can use VMware Aria OperationsVMware Cloud Foundation Operations to monitor and manage objects in your environment. These integrations are also referenced as Management Packs or as Solutions.

An integration can be a connection to a data source, or it might include predefined dashboards, widgets, alerts, and views.

Integrations can include cloud accounts, other accounts, dashboards, reports, alerts, and other content. The cloud accounts and other accounts comprise of adapters, using which VMware Aria OperationsVMware Cloud Foundation Operations manages communication and integration with other products, applications, and functions. When an integration is installed and the accounts are configured, you can use the VMware Aria OperationsVMware Cloud Foundation Operations analytics and alerting tools to manage the objects in your environment.

VMware Integrations include accounts for the following:
- Storage Devices
- VMware Aria Operations for Logs
- NSX for vSphere
- Network Devices
- VCM
- VMware Aria Hub

Third-party integrations include SCOM, EMC Smarts, and many others.

Other integrations such as the VMware Management Pack for NSX for vSphere, can be added to VMware Aria OperationsVMware Cloud Foundation Operations. To download VMware integrations and other third-party integrations, visit the VMware Marketplace.

VMware Aria OperationsVMware Cloud Foundation Operations includes integrations that are pre-installed.

VMware Aria OperationsVMware Cloud Foundation Operations also includes integrations that are bundled with VMware Aria OperationsVMware Cloud Foundation Operations, but not activated.

For a fresh deployment of VMware Aria OperationsVMware Cloud Foundation Operations, the activation status of integrations are as follows:

**Table 42: Integrations Activation Status**

| Integration Name | Activated by Default? | Can be Deactivated? |
|---|---|---|
| vCenter | Yes | No |
| VMware Cloud on AWS | No | Yes |
| vSAN | Yes | No |
| Service Discovery | Yes | No |
| VMware Aria Automation8.x | Yes | No |
| Azure VMware Solution | No | Yes |
| OS and Application Monitoring | Yes | No |
| Cloud Management Assessments | No | Yes |
| VMware Aria Operations for Logs | Yes | No |
| VMware Aria Operations for Networks | No | Yes |
| Google Cloud VMware Engine | No | Yes |
| VMware Cloud Foundation | Yes | Yes |
| NSX-T | Yes | No |
| VMware Aria Hub | No | Yes |
| Ping | No | Yes |
| PCI Compliance | No | Yes |
| ISO Compliance | No | Yes |
| HIPAA Compliance | No | Yes |
| FISMA Compliance | No | Yes |
| DISA Compliance | No | Yes |
| CIS Compliance | No | Yes |

The list of activated and available integrations, accounts, and repository are all available from a central **Integrations** page in VMware Aria OperationsVMware Cloud Foundation Operations. This page can be accessed from the left menu by clicking **Administration › Integrations**.

## Non-Native Management Packs and Management Pack Builder

If there's a technology without a current, native, management pack, you have the option to install a non-native pack or build your own using VMware Aria Operations Management Pack Builder.

VMware Aria Operations Management Pack Builder is a stand-alone appliance that enables the creation of custom management packs for use in VMware Aria Operations.  This is a no-code solution for bringing in data from an external API and either creating new resources or extending your VMware and third party resources with new Data, Relationships, and Events. For more information, see VMware Aria Operations for Integrations Index.

## Upgrade Considerations

The native integrations in VMware Aria OperationsVMware Cloud Foundation Operations are reinstalled if VMware Aria OperationsVMware Cloud Foundation Operations is upgraded.

If a new version of an integration is released, VMware sends an email to the organization owner with the new version information and the request to upgrade to the latest version. Integrations which require an update have the Upgrade Available sign on the top right corner of the integration card. With every release of a new version of an integration you must update to the latest version manually. If you fail to update to the new version by the time the next update for that integration is released, then VMware will send out a notice and update the pack to maintain compatibility since VMware supports the current version and one prior version on its cloud infrastructure.

# The Integrations Page in VMware Aria OperationsVMware Cloud Foundation Operations

The Integrations Page
Activate integrations, install integrations, add, import, or export accounts from the **Integrations** page in VMware Aria OperationsVMware Cloud Foundation Operations. Integrations are also referenced as Management Packs in VMware Aria OperationsVMware Cloud Foundation Operations.

### How Integrations Work

Integrations can include dashboards, reports, alerts, and accounts. Each account contains adapters using which VMware Aria OperationsVMware Cloud Foundation Operations integrates with other products, applications, and monitors their function and health.

### Where You Find Integrations

In the menu, click **Administration › Integrations**. By default, the **Accounts** tab opens. Use this tab to add, import, or export accounts. For more information, see The Accounts Tab.

Click the **Repository** tab to view the installed and available integrations. You can also install integrations delivered as PAK files. For more information, see The Repository Tab.

> **NOTE**
> For more information on pre-activated integrations, see Connecting to Data Sources.

### Data Collection Notifications

The **Data Collection** bell icon on the menu provides quick access to status and critical notifications related to data collections. The icon indicates whether notifications exist, and whether any of them are critical.

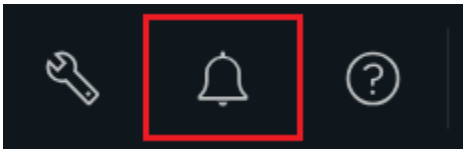

The list displays notifications about the data collections that are in progress, and indicates whether any of them have critical issues. The list groups the data collection notifications that are in progress into a single entry at the bottom of the list. To view the details about a collection, expand the notification.

Each notification displays the status of the last or current data collection, the associated adapter instance, and the time since the collection completed or an issue was identified. You can click a notification to open the Integrations page, where you can see further details, and manage adapter instances.

If problems occur with the data collections, VMware Aria OperationsVMware Cloud Foundation Operations identifies those problems during each 5-minute collection cycle.

### Failed Integration Installation

If a integration installation fails, plug-ins related to the integration might appear in the Plug-ins page of VMware Aria OperationsVMware Cloud Foundation Operations, even though the integration is not installed and does not appear on the Integrations page. When the integration installation fails, reinstall the integration.

## The Repository Tab

The repository tab lists all the installed and available integrations in VMware Aria OperationsVMware Cloud Foundation Operations. Integrations are also referenced as Management Packs and Solutions.

### Where You Find the Repository Tab

From the left menu, click **Administration › Integrations**. Click the **Repository** tab. The page displays installed integrations as clickable cards.

### What Can You Do in the Repository Tab

In the **Repository** tab, when you click on an integration card, the details page is displayed.

The management pack details page consists of the following three tabs:

- Overview. Describes the management pack and its purpose.
- Metrics. Displays a list of metrics and properties for the various object types as defined by the adapter.

  **NOTE**
  This list of metrics and properties can display additional metrics and properties after account creation since other sources can also publish them. The actual list of published metrics and properties depends from policy customizations and monitored environment.

- Content. Displays the various content types and the data defined by that management pack.

**Table 43: Management Pack Details Page Options**

| Options | Descriptions |
|---|---|
| Name | Name of the management pack, the name of the vendor or manufacturer who created the management pack, and version number. |
| Activate | Installs the management pack. You can configure cloud management packs after activation from the **Repository** or **Accounts** tab of the **Integrations** page.<br>The activation starts only if all the cluster's nodes are accessible.<br><br>**NOTE**<br>Some of the pre-Installed management packs are activated by default. You can configure them from the Management Pack Details Page or the **Accounts** tab of the **Integrations** page |
| Add Account | For more information on the accounts which are activated by default, see, Connecting to Data Sources. |
| **Horizontal Ellipses › Reset Default Content** | This option is only available for the vCenter management pack.<br>After you update your instance of VMware Aria OperationsVMware Cloud Foundation Operations and select the option |

*Table continued on next page*

*Continued from previous page*

| Options | Descriptions |
|---|---|
| | to overwrite, alert definitions and symptom definitions, you must overwrite your existing compliance alert definitions. |
| **Horizontal Ellipses › Deactivate** | You can uninstall a particular integration with its associated data, metadata, and the out of the box content. Select **I understand the risk and agree** to uninstall an integration. |

### Adding Solutions

Solutions are delivered as PAK files that you upload, license, and install.

#### How Added Solutions Work

When you add solutions, you configure adapters that manage the communication and integration between VMware Aria OperationsVMware Cloud Foundation Operations and other products, applications, and functionality.

#### Where You Add Solutions

From the left menu, click **Administration › Integrations**, and then click **Repository** in the right pane. Click **Add** to install other management packs. Click the vertical ellipsis of a management pack, and then click **Upgrade** to upgrade the management pack to the latest version.

From the left menu, click **Administration › Integrations**, and then click **Repository** in the right pane. Under **Available Integrations**, click **Get** to install the respective management pack. Any update available will appear under the **Upgrade Available** section. Click **Upgrade** to upgrade the management pack to the latest version.

#### Add Solutions Wizard Options

The wizard includes three pages where you locate and upload a PAK file, accept the EULA and install, and review the installation.

Before you install the PAK file, or upgrade your VMware Aria OperationsVMware Cloud Foundation Operations instance, clone any customized content to preserve it. Customized content can include alert definitions, symptom definitions, recommendations, and views.

While upgrading to the latest version, you can select the **Install the PAK file even if it is already installed** and the **Reset Default Content** options.

While upgrading to the latest version, you can select the **Reset Default Content** option.

**Table 44: Wizard Options**

| Option | Description |
|---|---|
| Page 1 | |
| Browse a Solution | Navigate to your copy of a management pack PAK file. |
| Download Solution | When you click **Get** in the **Repository** page, the solution is automatically downloaded. You can view the Name, Description, and Version of the Management Pack that is installed. |
| Upload | To prepare for installation, copy the PAK file to VMware Aria OperationsVMware Cloud Foundation Operations. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Install the PAK file even if it is already installed | If the PAK file was already uploaded, reload the PAK file using the current file, but leave user customizations in place. Do not overwrite or update the solution alerts, symptoms, recommendations, and policies. |
| Reset Default Content | If the PAK file was already uploaded, reload the PAK file using the current file, and overwrite the solution default alerts, symptoms, recommendations, and policies with newer versions provided with the current PAK file.<br>**NOTE**<br>A reset overwrites customized content. If you are upgrading VMware Aria OperationsVMware Cloud Foundation Operations, the best practice is to clone your customized content before you upgrade. |
| The PAK file is unsigned | Warning appears if the PAK file is not signed with a digital signature that VMware provides. The digital signature indicates the original developer or publisher and provides the authenticity of the management pack. If installing a PAK file from an untrusted source is a concern, check with the management pack distributor before proceeding with the installation. |
| Page 2 | |
| I accept the terms of the agreement | Read and agree to the end-user license agreement.<br>**NOTE**<br>Click **Next** to install the solution. The installation starts only if all the cluster's nodes are accessible. |
| Page 3 | |
| Installation Details | Review the installation progress, including the VMware Aria OperationsVMware Cloud Foundation Operations nodes where the adapter was installed. |

### The Accounts Tab

You can view and add accounts from the **Accounts** tab on the Integration page in VMware Aria OperationsVMware Cloud Foundation Operations. You can add accounts for pre-activated accounts or you can add accounts after activating integrations from the repository tab.

#### Where You Find the Accounts Tab

From the left menu, click **Administration › Integrations**. The **Accounts** tab opens by default. The accounts tab lists all the accounts configured for the different integrations available in VMware Aria OperationsVMware Cloud Foundation

Operations. You can filter them the accounts as per type to view accounts specific to an integration type.

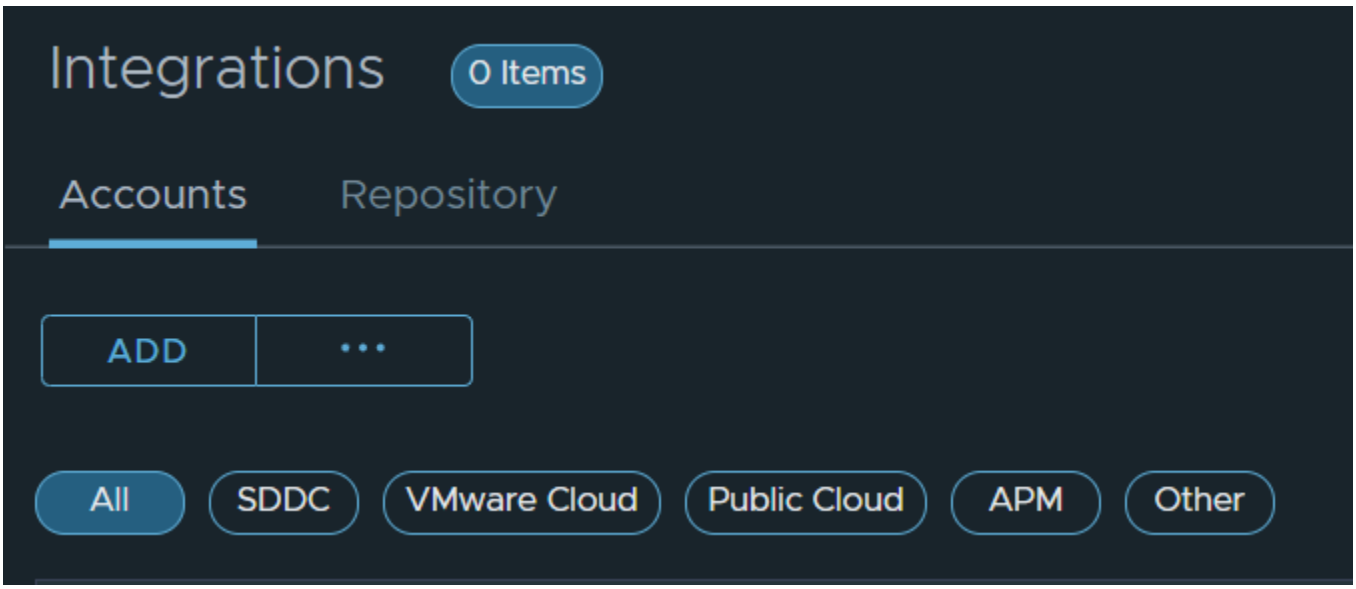


**What Can You Do in the Accounts Tab**

You can add and configure accounts associated with integrations that are pre-activated or those that you activate in VMware Aria OperationsVMware Cloud Foundation Operations. After you have configured the account, VMware Aria OperationsVMware Cloud Foundation Operations can collect data from or send data to the target system. You can access the accounts page at any time to modify your adapter configurations.

**Table 45: Accounts Page Options**

| Option | Description |
|---|---|
| Add | Click **Add** to configure accounts for pre-activated integrations. For more information on pre-activated integrations, see Connecting to Data Sources.<br>To add accounts for any of the other available integrations, select the account type and click **Yes** in the **Installation Required** wizard.<br>**NOTE**<br>You can also activate the integration from the Repository page and then add an account from the Accounts tab. For more information on activating available integrations from the Repository tab, see The Repository Tab. |
| Horizontal Ellipses | As a VMware Aria OperationsVMware Cloud Foundation Operations admin, you can backup the adapter configurations before upgrading. Click the horizontal ellipses and then click **Export Accounts** to export all the adapter configurations. To import them into a different VMware Aria OperationsVMware Cloud Foundation Operations instance, click **Import Accounts**. For more information, see Exporting and Importing Accounts. |

After configuring accounts you can view the account details and edit them in the Accounts page.

**Table 46: Accounts Grid Options**

| | |
|---|---|
| Vertical Ellipses | To view specific account details, expand the account type. Click the > icon next to the account and then click the vertical ellipses to change the account configurations.<br>• **Edit**. Allows you to edit the integrated adapter instance.<br>• **Delete**. Removes the adapter instance and clears the objects associated with the instance from the system, including historical data and role assignments.<br>**NOTE**<br>When you delete an account, you can choose to delete any related objects by selecting the checkbox, **Delete related objects**. If you do not wish to delete the related objects immediately, leave the check box unselected. The related objects are kept in the inventory for the duration of the retention period specified in the Global Settings page. If you recover an adapter instance before the end of the retention period, the related objects are unmarked and not deleted.<br>• **Delete All**. Available for adapter instances with multiple sub-adapter configurations (such as the vCenter adapter). It removes the adapter instance along with its child adapter instances.<br>**NOTE**<br>When you delete an account, you can choose to delete any related objects by selecting the checkbox, **Delete related objects**. If you do not wish to delete the related objects immediately, leave the check box unselected. The related objects are kept in the inventory for the duration of the retention period specified in the Global Settings page. If you recover an adapter instance before the end of the retention period, the related objects are unmarked and not deleted.<br>• **Start/ Stop Collecting**. Starts or Stops the data collection process.<br>• **Start/ Stop Collecting All**. Available for adapter instances with multiple sub-adapter configurations (such as the vCenter adapter). It starts or stops the data collection process of the adapter instance along with its child adapter instances.<br>• **Object details**. Open the object in the object browser. |
| Name | Name that the vendor or manufacturer gave to the solution. |

*Table continued on next page*

*Continued from previous page*

| | |
|---|---|
| Status | Indicates the status of the solution and whether the adapter is collecting any data. If the status displays a green tick with the text OK, it means that the solution is collecting data. |
| Description | Typically, an indication of what the solution monitors or what data source its adapter connects to. |
| Collector | Indicates the status of the solution. Data receiving shows that the solution is collecting data. |
| Filters | You can search the list of accounts according to the following criteria:<br>• Name<br>• Description<br>• Solution<br>• Adapter<br>• Collector |

**Adding Accounts**

You can add and configure accounts associated with integrations that are provided with or that you add to VMware Aria Operations. After you have configured the account, VMware Aria Operations can communicate with the target system. You can access the **Accounts** tab in the **Integrations** page at any time to modify your adapter configurations.

> **NOTE**
> Configure cloud proxy and activate the integration before adding and configuring accounts. For more information, see Configuring Cloud Proxies and The Repository Tab.

From the left menu, click **Administration › Integrations**. In the Accounts tab, click **Add** and select the integration you want to manage.
To manage accounts for the vSphere integration, see Cloud Account Information - Account Options.
To add and configure accounts for VMware Cloud on AWS, see Configuring VMware Cloud on AWS in
To add and configure accounts for Azure VMware Solution, see Configuring an Instance in .
To add and configure accounts for NSX-T, see Configuring the NSX-T Adapter.
To add and configure accounts for VMware Aria Log Insight, see Configuring VMware Aria Log Insight with VMware Aria Operations.
To add and configure accounts for VMware Aria Network Insight, see Configuring VMware Aria Network Insight.
To add and configure accounts for VMware Aria Automation 8.X, see Configuring VMware Aria Automation 8.x with VMware Aria Operations.

**Exporting and Importing Accounts**

As a VMware Cloud Foundation Operations admin, you can backup the adapter configurations before upgrading, export all the adapter configurations, and import them into a different VMware Cloud Foundation Operations instance. You can export the adapter configurations from VMware Cloud Foundation Operations on-prem to VMware Cloud Foundation Operations, VMware Cloud Foundation Operations to VMware Cloud Foundation Operations on-prem, VMware Cloud Foundation Operations on-prem to on-prem instance, and VMware Cloud Foundation Operations to cloud instance.

> **NOTE**
> Users with "Export" permission can export and users with "Import" permission can import the adapter configurations.

1. Export the adapter configuration.
   a) From the left menu, click **Administration › Integrations**.
   b) In the Accounts tab, select the adapter configuration(s) that you want to export, click the horizontal ellipses, and select **Export Accounts**.
   c) Setup a new password to export data. The password should be at least 14 characters long and must include at least one numerical character, upper and lower case character, and special character.
   d) Click **Export**.
   The adapter configurations are exported in .zip format. A password is used to encrypt the data. Use the same password while importing this file.
2. Import the adapter configuration.

   **NOTE**
   Before importing the content, ensure that you have exported the adapter configurations.

   a) From the left menu, click **Administration › Integrations**.
   b) Click the horizontal ellipses and select **Import Accounts**.
   c) Click **Browse** to select the `.zip` file and enter the password that you had set while exporting the content.
   d) If there is a conflict while importing the adapter configurations, you can either overwrite the existing adapter configurations or skip the import, which is the default option.
   e) Click **Import** to import adapter configurations to the destination setup.

### Configuring Credentials in Integrations

Credentials are the user accounts that VMware Aria OperationsVMware Cloud Foundation Operations uses to activate one or more integrations and associated adapters, and to establish communication with the target data sources. The credentials are supplied when you configure each adapter. You can create new credentials and modify the credential settings outside the adapter configuration or during adapter configuration to accommodate changes to your environment.

If you are modifying credentials to accommodate changes based on your password policy, you must validate your connection so that the adapters configured with these credentials begin using the new user name and password to communicate between VMware Aria OperationsVMware Cloud Foundation Operations and the target system.

Another use of credential management is to remove misconfigured credentials. If you delete valid credentials that were in active use by an adapter, you disable the communication between the two systems. You cannot delete an active credential that is being used by an adapter instance.

If you need to change the configured credential to accommodate changes in your environment, you can edit the credential settings without being required to configure a new adapter instance for the target system. To edit credential settings, from the left menu, click **Administration › Integrations**. In the Accounts tab, click **Credentials** .

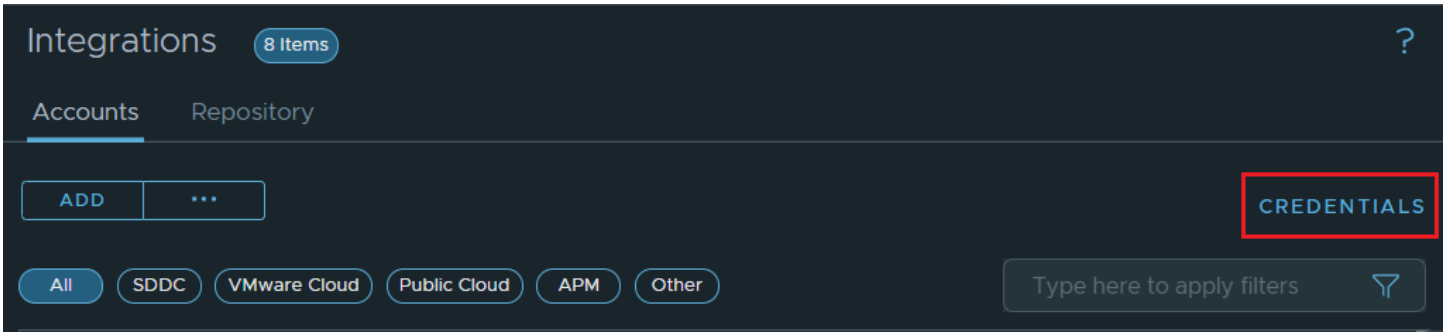

Starting VMware Aria OperationsVMware Cloud Foundation Operations 8.14, you could only view, add, modify, or delete credentials that you created or were assigned to you. You could view and use unassigned credentials only if you had the required permissions. When you upgrade to VMware Aria OperationsVMware Cloud Foundation Operations 8.16.1, you can deactivate the **Credential Ownership Enforcement** option from Global Settings to be able to modify credentials created and owned by others. For more information, see List of Global Settings.

Previously, you could only view, add, modify, or delete credentials that you created or that were assigned to you. You could view unassigned credentials only if you had the required permissions. When you upgrade to latest version of VMware Aria OperationsVMware Cloud Foundation Operations, you can deactivate the **Credential Ownership Enforcement** option from Global Settings to be able to modify credentials created and owned by others. For more information, see List of Global Settings.

**Where You Find Credentials**

From the left menu, click **Administration** › **Integrations**. In the **Accounts** tab, click the **Credentials** link on the upper right side.

**Table 47: Credentials Options**

| Option | Description |
|---|---|
| Toolbar options | Manages the selected credential.<br>• **Add**. Add new credentials for an adapter type that you can later apply when configuring an adapter.<br>• Click the **Vertical Ellipses** to perform any one of the following actions:<br>– **Edit**. Modify the selected credentials, usually when the user name and password require a change. The change is applied to the current adapter credentials after you validate the connection and the data source continues to communicate with VMware Aria OperationsVMware Cloud Foundation Operations.<br>– **Delete** . Remove the selected credentials from VMware Aria OperationsVMware Cloud Foundation Operations. If you have an adapter that uses these credentials, the communication fails and you cease monitoring the objects that the adapter was configured to manage. Commonly used to delete misconfigured credentials. You cannot delete an active credential that is being used by an adapter instance. |
| Filtering options | Limits the displayed credentials based on the adapter or credential types. |
| Credential name | Description of user-defined name that you provide to manage the credentials. Not the account user name. |
| Adapter Type | Adapter type for which the credentials are configured. |
| Credential Type | Type of credentials associated with the adapter. Some adapters support multiple types of credentials. For example, one type might define a user name and password, and another might define a pass code and key phrase. |

### Manage Credentials

To configure or reconfigure credentials that you use to activate an adapter instance, you must provide the collection configuration settings, for example, user name and password, that are valid on the target system. You can also modify the connection settings for an existing credential instance.

#### Where You Manage Credentials

From the left menu, click **Administration › Integrations**. In the **Accounts** tab, click the **Credentials** link on the upper right side.

#### Manage Credentials Options

The Manage Credentials dialog box is used to add new or modify existing adapter credentials. The dialog box varies depending on the type of adapter and whether you are adding or editing. The following options describe the basic options. Depending on the , the options other than the basic ones vary.

Starting VMware Aria OperationsVMware Cloud Foundation Operations 8.14, you could only view, add, modify, or delete credentials that you created or were assigned to you. You could view and use unassigned credentials only if you had the required permissions. When you upgrade to VMware Aria OperationsVMware Cloud Foundation Operations 8.16.1, you can deactivate the **Credential Ownership Enforcement** option from Global Settings to be able to modify credentials created and owned by others. For more information, see List of Global Settings.

Previously, you could only view, add, modify, or delete credentials that you created or that were assigned to you. You could view unassigned credentials only if you had the required permissions. When you upgrade to latest version of VMware Aria OperationsVMware Cloud Foundation Operations, you can deactivate the **Credential Ownership Enforcement** option from Global Settings to be able to modify credentials created and owned by others. For more information, see List of Global Settings.

**Table 48: Manage Credential Add or Edit Options**

| Option | Description |
|---|---|
| Adapter Type | Adapter type for which you are configuring the credentials. |
| Credential Kind | Credentials associated with the adapter. The combination of adapter and credential type affects the additional configuration options. |
| Credential Name | Descriptive name by which you are managing the credentials. |
| User Name | User account credentials that are used in the adapter configuration to connect VMware Aria OperationsVMware Cloud Foundation Operations to the target system. |
| Password | Password for the provided credentials. |

## vSphere

vSphere
vSphere
The vSphere solution connects VMware Aria OperationsVMware Cloud Foundation Operations to one or more vCenter instances. You collect data and metrics from those instances, monitor them, and run actions in them.

VMware Aria OperationsVMware Cloud Foundation Operations evaluates the data in your environment, identifying trends in object behavior, calculating possible problems and future capacity for objects in your system based on those trends, and alerting you when an object exhibits defined symptoms.

## Configuring the vSphere Solution

The vSphere solution is installed together with VMware Aria OperationsVMware Cloud Foundation Operations. The solution provides the vCenter adapter which you must configure to connect VMware Aria OperationsVMware Cloud Foundation Operations to your vCenter instances.

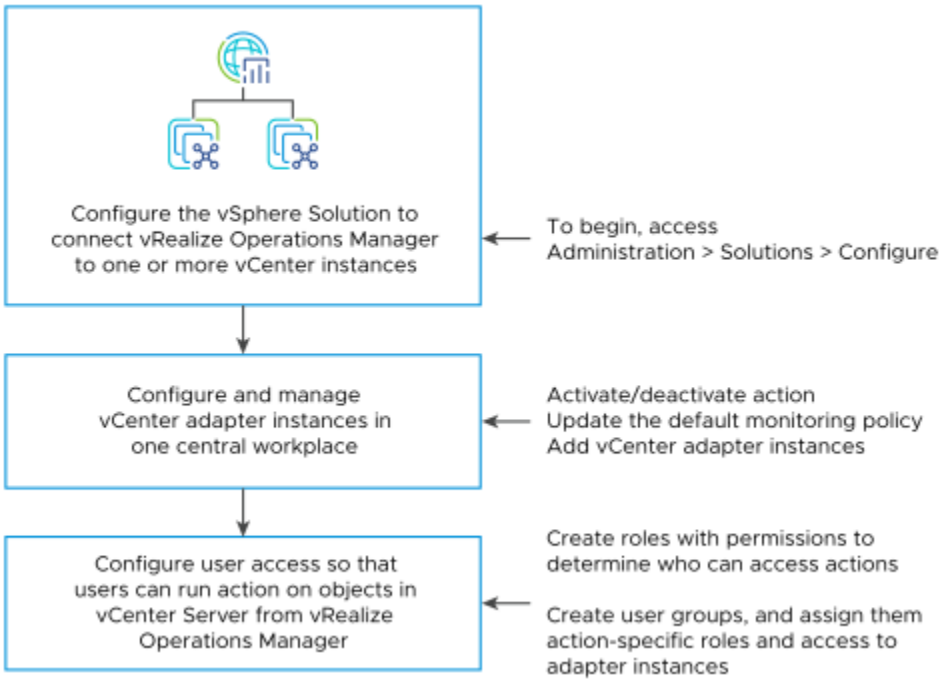


## How Adapter Credentials Work

The vCenter credentials that you use to connect VMware Aria OperationsVMware Cloud Foundation Operations to a vCenter instance, determines what objects VMware Aria OperationsVMware Cloud Foundation Operations monitors. Understand how these adapter credentials and user privileges interact to ensure that you configure adapters and users correctly, and to avoid some of the following issues.

- If you configure the adapter to connect to a vCenter instance with credentials that have permission to access only one of your three hosts, every user who logs in to VMware Aria OperationsVMware Cloud Foundation Operations sees only the one host, even when an individual user has privileges on all three of the hosts in the vCenter.
- If the provided credentials have limited access to objects in the vCenter, even VMware Aria OperationsVMware Cloud Foundation Operations administrative users can run actions only on the objects for which the vCenter credentials have permission.
- If the provided credentials have access to all the objects in the vCenter, any VMware Aria OperationsVMware Cloud Foundation Operations user who runs actions is using this account.

## Controlling User Access to Actions

Use the vCenter Server adapter to run actions on the vCenter Server from VMware Aria OperationsVMware Cloud Foundation Operations. If you choose to run actions, you must control user access to the objects in your vCenter Server environment. You control user access for local users based on how you configure user privileges in VMware Aria OperationsVMware Cloud Foundation Operations. If users log in using their vCenter account, then the way their account is configured in vCenter determines their privileges.

For example, you might have a vCenter user with a read-only role in vCenter. If you give this user the VMware Aria OperationsVMware Cloud Foundation Operations Power User role in vCenter rather than a more restrictive role, the user can run actions on objects because the adapter is configured with credentials that has privileges to change objects. To

avoid this type of unexpected result, configure local VMware Aria OperationsVMware Cloud Foundation Operations users and vCenter users with the privileges you want them to have in your environment.
To configure a vCenter cloud account, see Configure a Cloud Account in .

**Configure a vCenter Cloud Account in VMware Aria OperationsVMware Cloud Foundation Operations**

To manage your vCenter instances in VMware Aria OperationsVMware Cloud Foundation Operations, you must configure a cloud account for each vCenter instance. The cloud account requires the credentials that are used for communication with the target vCenter.

- Verify that you know the vCenter credentials that have sufficient privileges to connect and collect data, see Privileges Required for Configuring a vCenter Adapter Instance. If the provided credentials have limited access to objects in vCenter, all users, regardless of their vCenter privileges see only the objects that the provided credentials can access. At a minimum, the user account must have Read privileges and the Read privileges must be assigned at the data center or vCenter level.
- Ensure you configure your cloud proxy in VMware Cloud Foundation Operations, see Configuring Cloud Proxies in VMware Aria Operations.

  **NOTE**
  Any cloud account credentials you add are shared with other cloud account administrators and VMware Aria OperationsVMware Cloud Foundation Operations collector hosts. Other administrators might use these credentials to configure a new cloud account or to move a cloud account to a new host.

1. From the left menu, click **Administration › Integrations › Accounts** tab.
2. Click **Add**.
3. On the Accounts Type page, click **vCenter**.
4. Enter a display name and description for the cloud account.
   - Display name. Enter the name for the vCenter instance as you want it to appear in VMware Aria OperationsVMware Cloud Foundation Operations. A common practice is to include the IP address so that you can readily identify and differentiate between instances.
   - Description. Enter any additional information that helps you manage your instances.
5. Select the **Physical Data Center** you want to associate with the vCenter cloud account.

   If no physical data center is created or if you want to create a new physical data center for your cloud account, you can add a new physical data center. For more information, see Adding Physical Data Centers in
6. In the vCenter text box, enter the FQDN or IP address of the vCenter instance to which you are connecting.

   The vCenter FQDN or IP address must be reachable from all nodes in the VMware Aria OperationsVMware Cloud Foundation Operations cluster.

   The vCenter FQDN or IP address must be reachable from the selected cloud proxy.
7. To add credentials for the vCenter instance, click the **Add** icon, and enter the required credentials. The vCenter credential must have `Performance > Modify intervals` permission activated in the target vCenter to collect VM guest metrics.

   Optionally, you can use alternate user credentials for actions. Enter an **Action User Name** and **Password**. If you do not enter an action user name and password, the default user specified is considered for actions.

   **NOTE**
   Credentials are stored in VMware Aria OperationsVMware Cloud Foundation Operations and can be used for one or more instances of the vCenter.

**NOTE**
To monitor application services and operating systems, it is recommended that you enter action credentials with guest operations privileges such as
`guest operation alias modification`
,
`guest operation alias query`
,
`guest operation modifications`
,
`guest operation program execution`
,
`guest operation queries`
.

8. The collector for VMware Cloud Foundation Operations is the cloud proxy. Specify the cloud proxy you just deployed as the collector for this vCenter cloud account.
9. Determine which VMware Aria Operations collector or collector group is used to manage the cloud account. If you have only one cloud account, select **Default collector group**. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.
10. The collector for VMware Aria Operations can also be the cloud proxy. Select the cloud proxy you just deployed as the collector for this vCenter cloud account.
11. To collect data at an interval of 20 seconds from the vCenter, click **Activate** for Near Real-Time Monitoring. The Near Real-Time Monitoring dialog box appears. Click the **I understand the affects of near real-time monitoring** check box and then click **OK**.
12. The cloud account is configured to run actions on objects in the vCenter Server from VMware Aria Operations. If you do not want to run actions, deselect **Activate** for Operational Actions.
13. Click **Validate Connection** to validate the connection with your vCenter instance.
14. In the **Review and Accept Certificate** dialog box, review the certificate information.
    - If the certificate presented in the dialog box matches the certificate for your target vCenter, click **OK**.
    - If you do not recognize the certificate as valid, click **Cancel**. The test fails and the connection to vCenter is not completed. You must provide a valid vCenter URL or verify the certificate on the vCenter is valid before completing the adapter configuration.
15. To modify the advanced options regarding collectors, object discovery, or change events, expand the **Advanced Settings**.

    For information about these advanced settings, see Cloud Account Information - Account Options.

16. Click **Manage Registrations,** enter the credentials to register your VMware Aria OperationsVMware Cloud Foundation Operations instance in vCenter, and then click **Register**.
    - Enter username and password for a user with vCenter register/unregister privileges.
    - Select the collection credentials checkbox.
17. Click **Add** to save the configurations.

    The vCenter adapter instance gets saved.

The cloud account is added to the list. VMware Aria OperationsVMware Cloud Foundation Operations begins collecting metrics, properties, and events from the vCenter instance. Depending on the number of managed objects, the initial collection can take more than one collection cycle. A standard collection cycle begins every five minutes.

For information about the network port that VMware Aria OperationsVMware Cloud Foundation Operations uses to communicate with a vCenter system and VMware Aria OperationsVMware Cloud Foundation Operations components, see http://ports.vmware.com.

You can activate vSAN Configuration for your cloud account. For more information, see Configure a vSAN Adapter Instance .

You can use the vCenter for service discovery, see Configure Service and Application Discovery.

You can activate vSAN Configuration for your cloud account. For more information, see Configure a vSAN Adapter Instance.

You can use the vCenter for service discovery, see Configure Service Discovery.

### Privileges Required for Configuring a vCenter Adapter Instance

To configure your vCenter Adapter instance in VMware Aria OperationsVMware Cloud Foundation Operations, you need sufficient privileges to monitor and collect data and to perform vCenter actions. You can configure these permissions as a single role in vCenter to be used by a single service account or configure them as two independent roles for two separate service accounts.

The vCenter Adapter instance monitors and collects data from vCenter and the vCenter Action adapter performs some actions in vCenter. So, for monitoring or collecting vCenter inventory and their metrics and properties, the vCenter Adapter instance needs credentials with the following privileges activated in vCenter Server.

> **NOTE**
> The vCenter Server System Roles is created as a Read Only role with three system-defined privileges:: **System.Anonymous**, **System.View**, and **System.Read**. See, Using Roles to Assign Privileges.

**Table 49: Privileges for Configuring a vCenter Adapter: Monitoring and Data Collection**

| Task | Privilege |
|---|---|
| Property Collection | **System > Anonymous**<br>**NOTE**<br>This privilege is added automatically when you create a user account. However, this privilege is not visible in vSphere. |
| Objects Discovery<br>Events Collection | **Profile-Driven Storage > View**<br>**Storage views > View**<br>**Profile-Driven Storage > Profile-Driven Storage View**<br>**Datastore > Browse Datastore**<br>**System > View** |

*Table continued on next page*

*Continued from previous page*

| Task | Privilege |
|---|---|
| | **NOTE**<br>This privilege is added automatically when you create a user account. However, this privilege is not visible in vSphere. |
| Performance Metrics Collection | **Performance > Modify intervals**<br><br>**System > Read**<br><br>**NOTE**<br>This privilege is added automatically when you create a user account. However, this privilege is not visible in vSphere. |
| Service Discovery | For credential-based service discovery<br><br>**Virtual Machine > Guest Operations > Guest Operation alias modification**<br>**Virtual Machine > Guest Operations > Guest Operation alias query**<br><br>**Virtual Machine > Guest Operations > Guest Operation modifications**<br><br>**Virtual Machine > Guest Operations > Guest Operation program execution**<br><br>**Virtual Machine > Guest Operations > Guest Operation queries** |
| | For credential-less service discovery<br><br>**Virtual machine > Service configuration > Manage service configurations**<br>**Virtual machine > Service configuration > Modify service configuration**<br><br>**Virtual machine > Service configuration > Query service configurations**<br><br>**Virtual machine > Service configuration > Read service configuration** |
| VC Plugin | **Extension > Register extension**<br><br>**Extension > Unregister extension**<br><br>**Extension > Update extension** |
| Orphaned Disk | **Datastore > Browse datastore** |
| Authentication on VMware Aria OperationsVMware Cloud Foundation Operations using VC User and apply actions | **privilege.Global.com.vmware.label > VMware Aria Operations Read Only Role** |

*Table continued on next page*

*Continued from previous page*

| Task | Privilege |
|---|---|
| | **privilege.Global.com.vmware.label > VMware Aria Operations Power User Role** |
| Optimize Container<br><br>Schedule Optimize Container<br><br>Automate Optimize Container | • **AutoDeploy -> Rule -> Create**<br>• **AutoDeploy -> Rule -> Delete**<br>• **AutoDeploy -> Rule -> Edit**<br>• **AutoDeploy -> RuleSet -> Activate**<br>• **AutoDeploy -> RuleSet -> Edit**<br>• **Datastore -> Allocate Space**<br>• **Global -> Global tag**<br>• **Global -> System tag**<br>• **Host -> Inventory -> Manage Cluster Lifecyle**<br>• **Host -> Inventory -> Modify cluster**<br>• **Resource -> Assign virtual machine to resource pool**<br>• **Resource -> Migrate powered off virtual machine**<br>• **Resource -> Migrate powered on virtual machine**<br>• **Resource -> Query vMotion**<br>• **Storage views -> Configure service**<br>• **Storage views -> View**<br>• **Virtual machine -> Edit Inventory > Move**<br>Privilege required for vCenter version 7.x:<br>• **Profile-driven storage -> Profile-driven storage update**<br>• **Profile-driven storage -> Profile-driven storage view**<br><br>Privilege required for vCenter version 8.x :<br>• **VM storage policies -> Apply VM storage policies**<br>• **VM storage policies -> Update VM storage policies**<br>• **VM storage policies -> VM storage policies edit permissions**<br>• **VM storage policies -> VM storage policies view permissions**<br>• **VM storage policies -> View VM storage policies** |
| Provide data to vSphere Predictive DRS | **External stats provider > Update**<br><br>**External stats provider > Register**<br><br>**External stats provider > Unregister**<br><br>**vSphere Stats Privileges > Collect Stats Data**<br><br>**vSphere Stats Privileges > Modify Stats Configuration**<br><br>**vSphere Stats Privileges > Query Stats Data** |
| Tag Collection | **Global > Global tag**<br><br>**Global > Global health** |

*Table continued on next page*

*Continued from previous page*

| Task | Privilege |
|---|---|
| | **Global > Manage custom attributes**<br>**NOTE**<br>This privilege is required only if the tags are associated with custom attributes.<br>**Global > System tag**<br>**Global > Set custom attribute** |
| Monitoring and collecting data from vSphere with Tanzu | **Administrator**<br>**NOTE**<br>Users with Non-Administrator or custom role must be added to the **ServiceProviderUser** group. **Administrator > Single Sign On > Users and Groups > Groups.**<br>The ServiceProviderUsers is a group in the vCenter Single Sign-On Domain. Members of this group can manage the vSphere with Tanzu and VMware Cloud on AWS infrastructure. |
| Add License to vCenter | **Global. Licenses** |

**Table 50: Privileges for Configuring a vCenter Adapter: Performing vCenter Server Actions**

| Task | Privilege |
|---|---|
| Set CPU Count for VM | **Virtual Machine > Configuration > Change CPU Count** |
| Set CPU Resources for VM | **Virtual Machine > Configuration > Change Resource** |
| Set Memory for VM | **Virtual Machine > Configuration > Change Memory** |
| Set Memory Resources for VM | **Virtual Machine > Configuration > Change Resource** |
| Delete Idle VM | **Virtual machine > Edit Inventory > Remove** |
| Delete Powered Off VM | **Virtual machine > Edit Inventory > Remove** |
| Create Snapshot for VM | **Virtual Machine > Snapshot Management > Create Snapshot** |
| Delete Unused Snapshots for Datastore | **Virtual Machine > Snapshot Management > Remove Snapshot** |
| Delete Unused Snapshot for VM | **Virtual Machine > Snapshot Management > Remove Snapshot** |
| Power Off VM | **Virtual Machine > Interaction > Power Off** |
| Power On VM | **Virtual Machine > Interaction > Power On** |
| Shut Down Guest OS for VM | **Virtual Machine > Interaction > Power Off** |
| Move VM | • **Resource > Assign Virtual Machine to Resource Pool**<br>• **Resource > Migrate Powered Off Virtual Machine**<br>• **Resource > Migrate Powered On Virtual Machine**<br>• **Datastore > Allocate Space** |

*Table continued on next page*

*Continued from previous page*

| Task | Privilege |
|---|---|
| | • **Virtual machine -> Edit Inventory > Move**<br><br>**NOTE**<br>Combining these four permissions allows the service account to perform Storage vMotion and regular vMotion of an object therefore allowing VMware Aria OperationsVMware Cloud Foundation Operations to perform the given operations. |
| Set DRS Automation | **Host > Inventory > Modify Cluster** |
| Provide data to vSphere Predictive DRS | **External stats provider > Update**<br><br>**External stats provider > Register**<br><br>**External stats provider > Unregister** |
| Reboot Guest OS for VM | **Virtual machine > Interaction > Reset** |

For more information about tasks and privileges, see Required Privileges for Common Tasks in the *vSphere Virtual Machine Administration Guide* and Defined Privileges in the *vSphere Security Guide*.

**Configure User Access for Actions**

Configure User Access for Actions
Configure User Access for Actions
To ensure that users can run actions in VMware Aria OperationsVMware Cloud Foundation Operations, you must configure user access to the actions.

You use role permissions to control who can run actions. You can create multiple roles. Each role can give users permissions to run different subsets of actions. Users who hold the administrator role or the default super user role already have the required permissions to run actions.

You can create user groups to add action-specific roles to a group rather than configuring individual user privileges.

1. From the left menu, click **Control Panel**, and then click the **Access Control** tile.
2. To create a role:
   a) Click the **Roles** tab.
   b) Click the **Add** icon, and enter a name and description for the role.
3. To apply permissions to the role, select the role, and in the Permissions pane, click the **Edit** icon.
   a) Expand **Environment**, and then expand **Action**.
   b) Select one or more of the actions, and click **Update**.
4. To create a user group:
   a) Click the **User Groups** tab, and click the **Add** icon.
   b) Enter a name for the group and a description, and click **Next**.
   c) Assign users to the group, and click the **Objects** tab.
   d) Select a role that has been created with permissions to run actions, and select the **Assign this role to the user** check box.
   e) Configure the object privileges by selecting each adapter instance to which the group needs access to run actions.
   f) Click **Finish**.

Test the users that you assigned to the group. Log out, and log back in as one of the users. Verify that this user can run the expected actions on the selected adapter.

### Cloud Account Information - vSphere Account Options

To begin monitoring your environment with VMware Aria OperationsVMware Cloud Foundation Operations, you configure the vSphere solution. The solution includes the vCenter cloud account that collects data from the target vCenter instances.

#### Where You Find the Solution - vSphere

From the left menu, click **Administration › Integrations › Accounts** tab. Click **Add Account**, and then select the **vCenter** card.

#### Account Information - vSphere Account Options

Configure and modify cloud accounts on the Account Information page.

**Table 51: Advanced Settings Options**

| Option | Description |
|---|---|
| Advanced Settings | Provides options related to designating specific collectors to manage this cloud account, managing object discovery and change events. |
| Auto Discovery | Determines whether new objects added to the monitored system are discovered and added to VMware Aria OperationsVMware Cloud Foundation Operations after the initial configuration of the cloud account.<br>• If the value is true, VMware Aria OperationsVMware Cloud Foundation Operations collects the information about any new objects that are added to the monitored system after the initial configuration. For example, if you add more hosts and virtual machines, these objects are added during the next collections cycle. This is the default value.<br>• If the value is false, VMware Aria OperationsVMware Cloud Foundation Operations monitors only those objects that are present on the target system when you configure the cloud account. |
| Process Change Events | Determines whether the cloud account uses an event collector to collect and process the events generated in the vCenter instance.<br>• If the value is true, the event collector collects and publishes events from vCenter. This is the default value.<br>• If the value is false, the event collector does not collect and publish events. |
| Activate Collecting vSphere Distributed Switch<br><br>Activate Collecting Virtual Machine Folder<br><br>Activate Collecting vSphere Distributed Port Group | When set to false, reduces the collected data set by omitting collection of the associated category. |
| Exclude Virtual Machines from Capacity Calculations | When set to true, reduces the collected data set by omitting collection of the associated category. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Maximum Number Of Virtual Machines Collected | Reduces the collected data set by limiting the number of virtual machine collections.<br><br>To omit data on virtual machines and have VMware Aria OperationsVMware Cloud Foundation Operations collect only host data, set the value to zero. |
| Provide data to vSphere Predictive DRS | vSphere Predictive DRS proactively load balances a vCenter Server cluster to accommodate predictable patterns in the cluster workload.<br>VMware Aria OperationsVMware Cloud Foundation Operations monitors virtual machines running in a vCenter Server, analyzes longer-term historical data, and provides forecast data about predictable patterns of resource usage to Predictive DRS. Based on these predictable patterns, Predictive DRS moves to balance resource usage among virtual machines.<br><br>Predictive DRS must also be activated for the Compute Clusters managed by the vCenter Server instances monitored by VMware Aria OperationsVMware Cloud Foundation Operations. Refer to the *vSphere Resource Management Guide* for details on activating Predictive DRS on a per Compute Cluster basis.<br><br>When set to true, designates VMware Aria OperationsVMware Cloud Foundation Operations as a predictive data provider, and sends predicative data to the vCenter Server. You can only register a single active Predictive DRS data provider with a vCenter Server at a time. |
| Activate Actions | Activating this option helps in triggering the actions that are related to vCenter. |
| Cloud Type | Provides an ability to identify the type of vCenter that is used in VMware Aria OperationsVMware Cloud Foundation Operations. By default, the cloud type is set to Private Cloud.<br>The cloud types available are: Google Cloud VMware Engine, Hosted Private Cloud, Private Cloud, and VMware Cloud on AWS. |
| vCenter ID | A globally unique identifier associated with the vCenter instance. |
| Deactivate collecting Guest File Systems with names containing | Provide comma separated list of strings. If these strings are found in any guest files system mount point name, that guest file system will not be collected. |
| Activate real time monitoring | The real time monitoring setting is deactivated by default. To collect real time data every 20 seconds change it to **true**. |
| Dynamic Thresholding | This setting is activated by default. |

You can find the vSphere Hardening Guides at http://www.vmware.com/security/hardening-guides.html.

Click **Save Settings** to finish configuration of the solution.

## VMware Cloud on AWS

VMware Cloud on AWS provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing private cloud like operating environment. To manage your VMware Cloud on AWS instances in VMware Aria OperationsVMware Cloud Foundation Operations, you must configure a cloud account.

You can configure the following types of VMware Cloud on AWS endpoints in VMware Aria Operations:
- Commercial Cloud Endpoint
- Government Cloud Endpoint

### Configuring VMware Cloud on AWS in VMware Aria OperationsVMware Cloud Foundation Operations

To manage your VMware Cloud on AWS instances in VMware Aria OperationsVMware Cloud Foundation Operations, you must configure a cloud account. The adapter requires the CSP API token that is used to authorize and communicate with the target VMware Cloud on AWS.

Navigate to **API Tokens** under **My Account** and generate a CSP API token based on your operational needs:

- To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on AWS service roles.
- For data collection of bills, include either Billing Read-only or Organization Owner roles from All Organization Roles.

  **NOTE**
  The data collection of bills requires the bills to be available in the CSP.

- For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMware Cloud on AWS service roles.

1. From the left menu, click **Administration › Integrations**.
2. On the Accounts tab, click **Add**.
3. On the Accounts Types page, click **VMware Cloud on AWS**.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the VMware Cloud on AWS instance as you want it to appear in VMware Aria OperationsVMware Cloud Foundation Operations.
   - Description. Enter any additional information that helps you manage your instances.
5. To add credentials for the VMware Cloud on AWS instance, click the **Add** icon, and enter the required credentials.
   - Credential Name. The name by which you are identifying the configured credentials.
   - CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating a CSP API Token for VMware Cloud on AWS.

     **NOTE**
     Enter the following details if you are using proxy server to access internet or public services.

   - Proxy Host. A remote proxy server IP.
   - Proxy Port. The port that is activated on a remote proxy server.
   - Proxy username. Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as `username@domain name`.
   - Proxy Password. Password for the proxy server username.
   - Proxy Domain. The domain has to be empty while using the proxy with domain configuration.

     **NOTE**
     The proxy credentials will be used by NSX adapters.

6. Determine which VMware Cloud Foundation Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

   **NOTE**
   Ensure that you have Internet connectivity for the collectors to work.

7. Organization ID. Click **Get Organization** to auto-fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

   The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings › View Organization**.

8. Click **Validate Connection** to validate the connection.

9. You can monitor the costs of running your VMware Cloud on AWS infrastructure by bringing in the billing from VMware Cloud on AWS to VMware Aria OperationsVMware Cloud Foundation Operations. To do so, activate the costing option in **Advanced Settings**.

   **NOTE**
   If the bills are not available in CSP, then the VMC infrastructure costs calculation will automatically switch from the bill-based calculation to the list-price based calculation.

10. Click **Save**.
    The page to configure the SDDC in VMware Cloud on AWS appears.
11. From the list of available SDDCs in VMware Cloud on AWS, click any one of the SDDCs that you want to monitor from VMware Aria OperationsVMware Cloud Foundation Operations.
12. Configure the vCenter adapter:

    1. Click the **vCenter** tab, and enter the required credentials.
       - Credential Name. The name by which you are identifying the configured credentials.
       - User Name. The vCenter user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
       - Password. The vCenter password configured for that vCenter user name.
    2. Select the required collector group.

       **NOTE**
       If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of VMware Cloud on AWS.

       Select the Cloud Proxy deployed on the given VC and ensure it has access to the Internet. If the outbound internet access for the Cloud proxy must be restricted, ensure that the minimum Cloud Proxy prerequisites are met.

       For details, see "Configuring Cloud Proxies in VMware Cloud" topic in the *Getting Started with VMware Aria Operations* guide.

       It is advised not to use the default collector groups as the VMware Cloud on AWS management gateway firewall rule does not allow traffic originating from any address.

       If you have configured an HTTP proxy on your VMware Aria OperationsVMware Cloud Foundation Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

    3. If you have installed cloud proxy in VMware Cloud on AWS SDDC, the cloud proxy might not have outbound internet access to reach the VMware Aria OperationsVMware Cloud Foundation Operations service. To activate outbound internet access for the deployed cloud proxy and allow cloud proxy to connect to vCenter, perform the following steps:
       - Request a new public IP in the VMware Cloud on AWS SDDC where the cloud proxy was deployed. For details, see Request or Release a Public IP Address.
       - Add a new NAT rule for the internet that associates private IP of the cloud proxy with the public IP. For details, see Create or Modify NAT Rules.
       - Add a firewall rule that allows incoming traffic from the public IP that was associated with cloud proxy VM in earlier step to vCenter.

13. Click the **vSAN** tab. By default, the vSAN adapter is activated.

1. Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **Ok**.
2. Select **Enable SMART data collection**, if required.
3. Click **Validate Connection** to validate the connection.

14. Click the **NSX** tab. By default, the NSX adapter is activated.

    1. Click **Validate Connection** to validate the connection. If you have hardened the SDDC environment, then you may get an error while validating the NSX connection. To resolve this issue, change the NSX adapter instance to use the private IP address of NSX manager by performing the following steps:

        a. Navigate to **Operations › Configurations**, and then click the **Inventory Management** tile. Navigate to **Adapter Instances › NSX** and click your NSX adapter instance.
        b. From the list of objects displayed, edit the object of type NSX adapter instance and enter the private IP address of NSX manager for your environment in the **Virtual IP/NSX Manager** field.
        c. Click **OK**.

15. Click **Save This SDDC**.

    **NOTE**
    The Service Discovery adapter is optional. The steps to configure the VMware Cloud on AWS Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information, see Configure Service and Application Discovery.

The VMware Cloud on AWS account, with the configured SDDC, is added to the list.

### Known Limitations

Review the following list of feature limitations of VMware Cloud on AWS integration.

- The compliance workflows in VMware Cloud Foundation Operations work for the virtual machines running on a vCenter in Key definition for "cloud_serv_long" not found in the DITA map. on AWS. The compliance checks for VMware management objects such Hosts, vCenter, and so on, are not available.
- Workload optimization including pDRS and host-based business intent do not work because of the VMware managers cluster configurations.
- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with VMware Cloud Foundation OperationsVMware Cloud Foundation Operations. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter interface.
- Key definition for "cloud_serv_long" not found in the DITA map. does not support VMware Cloud Foundation Operations SaaS plugin.
- You cannot log in to VMware Cloud Foundation Operations using your Key definition for "cloud_serv_long" not found in the DITA map.vCenter credentials.
- Credential-less service discovery is not supported on Key definition for "cloud_serv_long" not found in the DITA map. on AWS.

### Generating a CSP API Token for VMware Cloud on AWS

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of Vmware Cloud on AWS.

- To configure the VMware Cloud on AWS Adapter, generate the CSP API token with any of the VMware Cloud on AWS service roles.
- For data collection of bills, generate the CSP API token with the Billing Read-only or Organization Owner organization role with any of the VMware Cloud on AWS service roles.
- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.
2. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.
4. Click **Generate**.
5. Copy or save the generated token.

### Verify that the NSX Adapter Instance is Connected and Collecting Data

You can verify if your adapter instance can retrieve information from the NSX objects in your inventory.

To view the object types, from the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile under Miscellaneous. Click **Adapter Instances › NSX**, and then click the user-created instance.

**Table 52: Object Types that NSX Discovers**

| Object Type | Description |
|---|---|
| NSX Adapter Instance | The VMware Aria OperationsVMware Cloud Foundation Operations Management Pack for NSX instance. |
| Logical Switch | Logical segments in the NSX environment. |
| Logical Switches | Group of the logical segments. |
| Firewall Section | Firewall sections in the NSX environment. |
| Firewall Sections | Group of firewall sections. |
| Logical Router | Logical routers in the NSX environment. |
| Logical Routers | Group of tier-0 and tier-1 logical routers. |
| Tier-0 Routers | Group of tier-0 logical routers. |
| Tier-1 Routers | Group of tier-1 logical routers. |
| Group | Groups in the NSX environment. |
| Management Groups | Group of management groups in the NSX environment. |
| Compute Groups | Group of compute groups in the NSX environment. |
| Groups | Group of both management and compute groups. |

1. In the menu, click **Operations › Configurations**, and then click the **Inventory Management** tile under Miscellaneous.
2. Click **Adapter Instances › NSX**, and then click the user-created instance.
3. Select the adapter instance name to display the list of objects discovered by your adapter instance.
4. Slide the display bar to the right to view the object status.

| Object Status | Description |
|---|---|
| Collection State | If green, the object is connected. |

*Table continued on next page*

*Continued from previous page*

| Object Status | Description |
|---|---|
| Collection Status | If green, the adapter is retrieving data from the object. |

5. Deselect the adapter instance name and expand the **Object Types** tag.

   Each Object Type name appears with the number of objects of that type in your environment.

### Configuring VMware Cloud on AWS Government Cloud Endpoint in VMware Cloud Foundation Operations

To manage your VMware Cloud on AWS instances in VMware Cloud Foundation Operations, you must configure a cloud account. The adapter requires the CSP API token that is used to authorize and communicate with the target VMware Cloud on AWS.

- Navigate to **API Tokens** under **My Account** and generate a CSP API token based on your operational needs:
  - To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on AWS service roles.
  - For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMware Cloud on AWS service roles.
- To activate the cost calculations based on VMware Cloud on AWS GovCloud Pricing, you must modify the VMware Cloud on AWS rate card on the Cloud Providers tab in the Cost Settings page. For details on updating the rate card, see the VMware KB article 88488.

1. From the left menu, click **Administration › Integrations**.
2. On the Accounts tab, click **Add**.
3. On the Account Types page, click **VMware Cloud on AWS**.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the VMware Cloud on AWS instance as you want it to appear in VMware Cloud Foundation Operations.
   - Description. Enter any additional information that helps you manage your instances.
5. Under Advanced Settings, set **Environment** to **VMware Government Cloud** and **Billing Enabled** to **False**.
6. To add credentials for the VMware Cloud on AWS instance, click the **Add** icon, and enter the required credentials.
   - Credential Name. The name by which you are identifying the configured credentials.
   - CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating CSP API Token.

     **NOTE**
     Enter the following details if you are using a proxy server to access the Internet or public services.
   - Proxy Host. A remote proxy server IP.
   - Proxy Port. The port that is activated on a remote proxy server.
   - Proxy username. Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as `username@domain name`.
   - Proxy Password. Password for the proxy server username.
   - Proxy Domain. The domain has to be empty while using the proxy with domain configuration.

     **NOTE**
     The proxy credentials will be used by NSX adapters.
7. Determine which VMware Cloud Foundation Operations collector or collector group is used in managing the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the

workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**NOTE**
It is recommended to use a dedicated Cloud Proxy in VMC on AWS government cloud to monitor the VMC on AWS government cloud endpoint.

Ensure that you have Internet connectivity for the collectors to work.

If you have installed cloud proxy in VMware Cloud on AWS government cloud SDDC, the cloud proxy might not have outbound internet access to reach the VMware Cloud Foundation Operations service. To activate outbound internet access for the deployed cloud proxy and allow cloud proxy to connect to vCenter, perform the following steps:

a) Request a new public IP in the VMware Cloud on AWS government cloud SDDC where the cloud proxy was deployed. For details, see Request or Release a Public IP Address.
b) Add a new NAT rule for the internet that associates private IP of the cloud proxy with the public IP. For details, see Create or Modify NAT Rules.
c) Add a firewall rule that allows incoming traffic from the public IP that was associated with cloud proxy VM in earlier step to vCenter.

8. Organization ID. Click **Get Organization** to auto-fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

   The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings › View Organization**.

9. Under **Advanced Settings**, set **Billing Enabled** to **False**.
10. Click **Validate Connection** to validate the connection.
11. Click **Save**.
    The page to configure the SDDC in VMware Cloud on AWS appears.
12. From the list of available SDDCs in VMware Cloud on AWS government cloud, click any one of the SDDCs that you want to monitor from VMware Cloud Foundation Operations.
13. Configure the vCenter adapter:

    1. Click the **vCenter** tab, and enter the required credentials.
       – Credential Name. The name by which you are identifying the configured credentials.
       – User Name. The vCenter user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
       – Password. The vCenter password configured for that vCenter user name.
    2. Select the required collector group.

**NOTE**
It is recommended to use a dedicated Cloud Proxy in VMC on AWS government cloud to monitor the VMC on AWS government cloud endpoint.

If you are using a private IP for your vCenter Server or if you want to deploy telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of VMware Cloud on AWS.

Select the Cloud Proxy deployed on the given VC and ensure it has access to the Internet. If the outbound internet access for the Cloud proxy must be restricted, ensure that the minimum Cloud Proxy prerequisites are met. For details, see "Configuring Cloud Proxies in VMware Cloud" topic in the *Getting Started with VMware Aria Operations* guide.

It is advised not to use the default collector groups as the VMware Cloud on AWS management gateway firewall rule does not allow traffic originating from any address.

If you have configured an HTTP proxy on your VMware Cloud Foundation Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

3. If you have installed cloud proxy in VMware Cloud on AWS government cloud SDDC, the cloud proxy might not have outbound internet access to reach the VMware Cloud Foundation Operations service. To activate outbound internet access for the deployed cloud proxy and allow cloud proxy to connect to vCenter, perform the following steps:
   - Request a new public IP in the VMware Cloud on AWS government cloud SDDC where the cloud proxy was deployed. For details, see Request or Release a Public IP Address.
   - Add a new NAT rule for the internet that associates private IP of the cloud proxy with the public IP. For details, see Create or Modify NAT Rules.
   - Add a firewall rule that allows incoming traffic from the public IP that was associated with cloud proxy VM in earlier step to vCenter.
   - Click **Next**.

14. Configure the vSAN adapter.

    1. Click the **vSAN** tab. By default, the vSAN adapter is activated.
    2. Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **Ok**.
    3. Select **Enable SMART data collection**, if required.
    4. Click **Validate Connection** to validate the connection.
    5. Click **Next**.

15. Configure the NSX adapter.

    1. Click the **NSX** tab.
    2. Activate NSX configuration if it is deactivated.
    3. Click **Validate Connection** to validate the connection. If you have hardened the SDDC environment, then you may get an error while validating the NSX connection. To resolve this issue, change the NSX adapter instance to use the private IP address of NSX manager by performing the following steps:

       a. Navigate to **Operations › Configurations**, and then click the **Inventory Management** tile. Navigate to **Adapter Instances › NSX** and click your NSX adapter instance.

b. From the list of objects displayed, edit the object of type NSX adapter instance and enter the private IP address of NSX manager for your environment in the **Virtual IP/NSX Manager** field.
c. Click **Next**.

16. Click **Save This SDDC**.

**NOTE**
The Service Discovery adapter is optional. The steps to configure the VMware Cloud on AWS Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information, see Configure Service and Application Discovery.

The VMware Cloud on AWS government cloud account, with the configured SDDC, is added to the list.

### Known Limitations

- For Government Cloud Endpoint, bill-based costing is not supported and the costing defaults to rate card based costing.
- The compliance workflows in VMware Cloud Foundation Operations checks for compliance only for virtual machines running on a vCenter servcer in VMware Cloud on AWS. The compliance checks for VMware management objects such as Hosts, vCenter, and so on, are not available.
- Workload optimization including pDRS and host-based business intent does not work because VMware manages cluster configurations.
- Workload optimization for the cross-cluster placement within the SDDC with the cluster-based business intent is fully supported with VMware Aria OperationsVMware Cloud Foundation Operations. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter interface.
- Key definition for "cloud_serv_long" not found in the DITA map. does not support the VMware Cloud Foundation Operations SaaS plugin.
- You cannot log in to VMware Cloud Foundation Operations using your Key definition for "cloud_serv_long" not found in the DITA map.vCenter credentials.
- Credential-less service discovery is not supported on Key definition for "cloud_serv_long" not found in the DITA map. on AWS.

### Generating CSP API Token

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of VMware Cloud on AWS.

- To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMware Cloud on AWS service roles.
- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.
2. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.
4. Click **Generate**.
5. Copy or save the generated token.

### Verify that the NSX Adapter Instance is Connected and Collecting Data

You can verify if your adapter instance can retrieve information from the NSX objects in your inventory.

To view the object types, from the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile under Miscellaneous. Click **Adapter Instances › NSX**, and then click the user-created instance.

**Table 53: Object Types that NSX Discovers**

| Object Type | Description |
|---|---|
| NSX Adapter Instance | The VMware Aria OperationsVMware Cloud Foundation Operations Management Pack for NSX instance. |
| Logical Switch | Logical segments in the NSX environment. |
| Logical Switches | Group of the logical segments. |
| Firewall Section | Firewall sections in the NSX environment. |
| Firewall Sections | Group of firewall sections. |
| Logical Router | Logical routers in the NSX environment. |
| Logical Routers | Group of tier-0 and tier-1 logical routers. |
| Tier-0 Routers | Group of tier-0 logical routers. |
| Tier-1 Routers | Group of tier-1 logical routers. |
| Group | Groups in the NSX environment. |
| Management Groups | Group of management groups in the NSX environment. |
| Compute Groups | Group of compute groups in the NSX environment. |
| Groups | Group of both management and compute groups. |

1. In the menu, click **Operations › Configurations**, and then click the **Inventory Management** tile under Miscellaneous.
2. Click **Adapter Instances › NSX**, and then click the user-created instance.
3. Select the adapter instance name to display the list of objects discovered by your adapter instance.
4. Slide the display bar to the right to view the object status.

| Object Status | Description |
|---|---|
| Collection State | If green, the object is connected. |
| Collection Status | If green, the adapter is retrieving data from the object. |

5. Deselect the adapter instance name and expand the **Object Types** tag.

   Each Object Type name appears with the number of objects of that type in your environment.

## VMware Cloud on AWS Outposts

VMware Cloud on AWS Outposts provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing private cloud like operating environment. To manage your VMware Cloud on AWS Outposts instances in VMware Aria OperationsVMware Cloud Foundation Operations, you must configure a cloud account.

### Configuring VMware Cloud on AWS Outposts in VMware Aria OperationsVMware Cloud Foundation Operations

To manage your VMware Cloud on AWS Outposts instances in VMware Aria OperationsVMware Cloud Foundation Operations, you must configure a cloud account. The adapter requires the CSP API token that is used to authorize and communicate with the target VMware Cloud on AWS Outposts.

Navigate to **API Tokens** under **My Account** and generate a CSP API token based on your operational needs:

- To discover and manage SDDCs, include Administrator (Delete Restricted) or Administrator from VMWare Cloud on AWS service roles.
- For data collection of bills, include either Billing Read-only or Organization Owner roles from All Organization Roles.

  **NOTE**
  The data collection of bills requires the bills to be available in the CSP.

- For NSX monitoring, include NSX Cloud Admin or NSX Cloud Auditor roles from VMWare Cloud on AWS service roles.
- To activate the cost calculations based on VMware Cloud on AWS Outposts pricing, you must modify the VMware Cloud on AWS Outposts rate card on the Cloud Providers tab in the Cost Settings page. For details on updating the rate card, see the VMware KB article KB88488.
- If you have subscribed to both VMC on AWS service and VMware Cloud on AWS Outposts service, ensure that they are in different CSP organizations. Otherwise, certain functionalities such as costing, and configuration maximums may not function as expected when both VMC on AWS service and VMware Cloud on AWS Outposts services are onboarded for monitoring within VMware Aria OperationsVMware Cloud Foundation Operations.

1. From the left menu, click **Administration › Integrations**.
2. On the Accounts tab, click **Add**.
3. On the Accounts Types page, click **VMware Cloud on AWS**.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the VMware Cloud on AWS Outposts instance as you want it to appear in VMware Aria OperationsVMware Cloud Foundation Operations.
   - Description. Enter any additional information that helps you manage your instances.
5. To add credentials for the VMware Cloud on AWS Outposts instance, click the **Add** icon, and enter the required credentials.
   - Credential Name. The name by which you are identifying the configured credentials.
   - CSP Refresh Token. A CSP API token. For details on generating an API token, see Generating a CSP API Token for .

     **NOTE**
     Enter the following details if you are using proxy server to access internet or public services.

   - Proxy Host. A remote proxy server IP.
   - Proxy Port. The port that is available on a remote proxy server.
   - Proxy username. Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as `username@domain name`.
   - Proxy Password. Password for the proxy server username.
   - Proxy Domain. The domain has to be empty while using the proxy with domain configuration.

     **NOTE**
     The proxy credentials will be used by NSX adapters.

6. Determine which VMware Aria OperationsVMware Cloud Foundation Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

**NOTE**
The CSP token is used to access the publicly available VMware Cloud Services Portal API. It is recommended to use the **Default Collector Group** for this access. If you use a Cloud Proxy, ensure it has access to the Internet, or if the outbound internet access for the Cloud Proxy must be restricted, ensure the minimum Cloud Proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Aria Operations.

If you have installed Cloud Proxy in VMware Cloud on AWS Outposts SDDC, the Cloud Proxy might not have outbound internet access to reach the VMware Cloud Foundation Operations service. To activate outbound internet access for the deployed Cloud Proxy and allow Cloud Proxy to connect to vCenter, perform the following steps:

- Request a new public IP in the VMware Cloud on AWS Outposts SDDC where the Cloud Proxy was deployed. For more details, see Request or Release a Public IP Address.
- Add a new NAT rule for the internet that associates private IP of the Cloud Proxy with the public IP. For details, see Create or Modify NAT Rules.
- Add a firewall rule that allows incoming traffic from the public IP that was associated with Cloud Proxy VM in earlier step to vCenter.

**NOTE**
Ensure that you have Internet connectivity for the collectors to work.

7. Organization ID. Click **Get Organization** to auto-fill this field. If you are offline or if you are unable to get the Organization ID, you can enter it manually.

   The Organization ID refers to the Long Organization ID in the Cloud Service Portal. To obtain this ID in the Cloud Service Portal, click **Organization Settings › View Organization**.

8. Click **Validate Connection** to validate the connection.
9. Bill-based costing is not supported in VMware Cloud on AWS Outposts. In the advanced settings, set "Billing Enabled" field to false
10. Click **Save**.
    The page to configure the SDDC in VMware Cloud on AWS Outposts appears.
11. From the list of available SDDCs in VMware Cloud on AWS Outposts, click any one of the SDDCs that you want to monitor from VMware Aria OperationsVMware Cloud Foundation Operations.
12. Configure the vCenter adapter:

    1. Click the **vCenter** tab, and enter the required credentials.
       – Credential Name. The name by which you are identifying the configured credentials.
       – User Name. The vCenter user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
       – Password. The vCenter password configured for that vCenter user name.
    2. Select the required collector group.

**NOTE**
If you have direct connectivity with your VMware Cloud vCenter Server, select **Default collector group**. If you are using a private IP for your vCenter Server or if you want to deploy telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of VMware Cloud on AWS Outposts.

Select the Cloud Proxy deployed on the given VC and ensure it has access to the Internet. If the outbound internet access for the Cloud proxy must be restricted, ensure that the minimum Cloud Proxy prerequisites are met.

For details, see the "Configuring Cloud Proxies in VMware Aria Operations" topic in VMware Cloud Foundation Operations on-prem *Getting Started* guide.

For details, see Configuring Cloud Proxies in VMware Aria Operations.

It is advised not to use the default collector groups as the VMware Cloud on AWS Outposts management gateway firewall rule does not allow traffic originating from any address.

If you have configured an HTTP proxy on your VMware Cloud Foundation Operations Cloud Proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

3. If you have installed a Cloud Proxy in VMware Cloud on AWS Outposts SDDC, the Cloud Proxy might not have outbound internet access to reach the VMware Cloud Foundation OperationsVMware Aria Operations service. To activate outbound internet access for the deployed Cloud Proxy and allow Cloud Proxy to connect to vCenter, perform the following steps:
   – Request a new public IP in the VMware Cloud on AWS Outposts SDDC where the Cloud Proxy was deployed. For details, see Request or Release a Public IP Address.
   – Add a new NAT rule for the internet that associates private IP of the Cloud Proxy with the public IP. For details, see Create or Modify NAT Rules.
   – Add a firewall rule that allows incoming traffic from the public IP that was associated with Cloud Proxy VM in earlier step to vCenter.

13. Click the **vSAN** tab. By default, the vSAN adapter is activated.

    1. Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **Ok**.
    2. Select **Enable SMART data collection**, if required.
    3. Click **Validate Connection** to validate the connection.

14. Click the **NSX** tab. By default, the NSX adapter is activated.

    1. Click **Validate Connection** to validate the connection. If you have hardened the SDDC environment, then you may get an error while validating the NSX connection. To resolve this issue, change NSX adapter instance to use the private IP address of NSX manager by following the steps below:

       a. Navigate to **Operations › Configurations**, and then click the **Inventory Management** tile. Navigate to **Adapter Instances › NSX** and click your NSX adapter instance.
       b. In the list of objects shown, edit the object of type NSX adapter instance and enter the private IP address of NSX manager for your environment in the **Virtual IP/NSX Manager** field.
       c. Click **OK**.

15. Click **Save This SDDC**.

**NOTE**
The Service Discovery adapter is optional. The steps to configure the VMware Cloud on AWS Outposts Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information about configuring the vCenter Service Discovery. see *Configure Service Discovery*.

The VMware Cloud on AWS Outposts account, with the configured SDDC, is added to the list.

### Known Limitations

Review the following list of feature limitations of VMware Cloud on AWS Outposts integration.

- VMware Cloud on AWS Outposts does not support Bill Based Costing.
- VMware Cloud on AWS Outposts does not support configuration maximums. You can ignore the displayed configuration maximums and related alerts.
- The compliance workflows in VMware Cloud Foundation Operations work for the virtual machines running on a vCenter in VMware Cloud on AWS Outposts. The compliance checks for VMware management objects such Hosts, vCenter, and so on, are not available.
- Workload optimization including pDRS and host-based business intent do not work because of the VMware Cloud Foundation OperationsVMware Cloud Foundation Operations cluster configurations in VMware Cloud on AWS Outposts.
- Workload optimization for the cross cluster placement within the SDDC with the cluster-based business intent is fully supported with VMware Cloud Foundation OperationsVMware Cloud Foundation Operations. However, workload optimization is not aware of resource pools and places the virtual machines at the cluster level. A user can manually correct this in the vCenter interface.
- VMware Cloud on AWS Outposts does not support VMware Cloud Foundation Operations plugin.
- You cannot log in to VMware Cloud Foundation Operations using your vCenter credentials.
- Credential-less service discovery is not supported on VMware Cloud on AWS Outposts.
- If you have subscribed to both VMC on AWS service and VMware Cloud on AWS Outposts service, ensure that they are in different CSP organizations. Otherwise, certain functionalities such as costing, and configuration maximums may not function as expected when both VMC on AWS service and VMware Cloud on AWS Outposts services are onboarded for monitoring within VMware Aria OperationsVMware Cloud Foundation Operations.

### Generating a CSP API Token for VMware Cloud on AWS Outposts

After a user is onboarded to the VMware Cloud Services, an account is created for that user. The user can log in to the account and generate an API token that can be configured as part of VMware Cloud on AWS Outposts.

- To configure the VMware Cloud on AWS Outposts Adapter, generate the CSP API token with any of the VMware Cloud on AWS service roles.
- For data collection of bills, generate the CSP API token with the Billing Read-only or Organization Owner organization role with any of the VMware Cloud on AWS service roles.
- For NSX monitoring, generate the CSP API token with the NSX Cloud Admin or NSX Cloud Auditor VMware Cloud on AWS service role.

1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.
2. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.
3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.
4. Click **Generate**.
5. Copy or save the generated token.

### Verify that the NSX Adapter Instance is Connected and Collecting Data

You can verify if your adapter instance can retrieve information from the NSX objects in your inventory.

To view the object types, from the left menu, click **Inventory › NSX**, and then click the user-created instance.

**Table 54: Object Types that NSX Discovers**

| Object Type | Description |
|---|---|
| NSX Adapter Instance | The Aria Operations management pack for the NSX instance. |
| Logical Switch | Logical segments in the NSX environment. |
| Logical Switches | Group of the logical segments. |
| Firewall Section | Firewall sections in the NSX environment. |
| Firewall Sections | Group of firewall sections. |
| Logical Router | Logical routers in the NSX environment. |
| Logical Routers | Group of tier-0 and tier-1 logical routers. |
| Tier-0 Routers | Group of tier-0 logical routers. |
| Tier-1 Routers | Group of tier-1 logical routers. |
| Group | Groups in the NSX environment. |
| Management Groups | Group of management groups in the NSX environment. |
| Compute Groups | Group of compute groups in the NSX environment. |
| Groups | Group of both management and compute groups. |

1. In the menu, click **Inventory**, and then click **NSX**.
2. Select the adapter instance name to display the list of objects discovered by your adapter instance.
3. Go to the **Details** tab.
4. Slide the display bar to the right to view the object status.

| Object Status | Description |
|---|---|
| Collection State | If green, the object is connected. |
| Collection Status | If green, the adapter is retrieving data from the object. |

5. Deselect the adapter instance name and expand the **Object Types** tag.

   Each Object Type name appears with the number of objects of that type in your environment.

## Azure VMware Solution

Azure VMware Solution provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing a private cloud like operating environment.

### Configuring an Azure VMware Solution Instance in VMware Aria OperationsVMware Cloud Foundation Operations

To monitor Azure VMware Solution instances in VMware Aria OperationsVMware Cloud Foundation Operations, you must add an Azure VMware Solution cloud account.

1. From the left menu, click **Administration › Integrations**.
2. In the Accounts tab, click **Add**.
3. On the Accounts Type page, click **Azure VMware Solution**.
4. From the **Add Cloud Account** page, enter a display name and description for the cloud account.

- Name. Enter the name for the Azure VMware Solution instance as you want it to appear in VMware Cloud Foundation OperationsVMware Cloud Foundation Operations.
- Description. Enter any additional information that helps you manage your instances.

5. Configure the Azure VMware Solution credentials.

| Option | Description |
|---|---|
| Subscription ID | Enter your subscription ID for Microsoft Azure. |
| Directory (Tenant) ID | Enter the directory (tenant) ID for your Azure Active Directory |
| Credential | Add the credentials used to access Azure VMware Solution by clicking the plus sign.<br>• Enter an instance name for the credential values you are creating. This value is not the name of the adapter instance, but a friendly name for the secret credential.<br>• Enter your application (Client) ID in your Azure Active Directory.<br>• Enter the client secret that you generated for your application in the Microsoft Azure portal.<br>**NOTE**<br>For steps to generate a client secret, see Generate a Client Secret for Azure VMware Solution.<br>**NOTE**<br>The adapter uses the Azure public APIs to fetch the private clouds, hence internet connectivity to poll the APIs listed under Private Clouds - List in Subscription is mandatory. To view the APIs that the adapter uses internally, see Resource Manager - Overview.<br>• Enter any required local proxy information for your network. |
| Collector/Group | Determine which VMware Cloud Foundation Operations collector or collector group is used to manage the cloud account.<br>The best practice is to deploy the cloud proxy on each Private Cloud instance of Azure VMware Solution. If you use a cloud proxy, ensure it has access to the Internet, or if the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Aria OperationsConfiguring Cloud Proxies in VMware Aria Operations. |

6. Click **Validate Connection** to validate the connection.
7. Click **Save**.
   The page to configure the Azure VMware Solution as a private cloud appears.
8. Click **Configure**.
9. Configure the desired vCenter adapter instance:

1. Click the **Add** icon against **Credential**, and enter the required credentials.
   - Credential Name. The name by which you are identifying the configured credentials.
   - User Name. The vCenter user name.
   - Password. The vCenter password configured for that vCenter user name.
2. Select the required collector group.

   **NOTE**
   If you have direct connectivity with your VMware Cloud vCenter, select **Default collector group**. If you are using a private IP for your vCenter or if you want to deploy Telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each SDDC instance of Azure VMware Solution.

   Select the cloud proxy deployed on the given vCenter and ensure it has access to the Internet. If the outbound internet access for the cloud proxy must be restricted, ensure that the minimum cloud proxy prerequisites are met.

   For details, see Configuring Cloud Proxies in VMware Aria Operations.

   For details, see Configuring Cloud Proxies in VMware Aria Operations.

   It is advised not to use the default collector groups as the Azure VMware Solution management gateway firewall rule does not allow traffic originating from any address.

   If you have configured an HTTP proxy on your VMware Aria OperationsVMware Cloud Foundation Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.
3. Click **Next** to navigate to the **vSAN** section.

10. By default, the vSAN adapter is activated.

    1. By default, the vCenter referenced credential will be used for vSAN validation.
    2. Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **OK**.
    3. Select **Enable SMART data collection**, if required.
    4. Click **Validate Connection** to validate the connection.
    5. Click **Next**.

11. Configure the NSX adapter.

    1. By default, the NSX configuration is activated.
    2. Click the **Add** icon against **Credential**, and enter the required credentials.
       - Credential Kind: Select the configured NSX instance.
       - Credential Name. The name by which you are identifying the configured credentials.
       - User Name. The user name of the NSX instance.
       - Password. The password of the NSX instance.
       - Click **OK**.
    3. Click **Validate Connection** to validate the connection.

12. Click **Save This Private Cloud**.

**NOTE**
The Service Discovery adapter is optional. The steps to configure the Azure VMware Solution Service Discovery adapter are similar to configuring vCenter Service Discovery. For more information about configuring the vCenter Service Discovery, see Configure Service and Application Discovery.

13. Determine which VMware Cloud Foundation Operations collector or collector group is used to manage the cloud account.

    The best practice is to deploy the cloud proxy on each private cloud instance of Azure VMware Solution. If you use a cloud proxy, ensure that it has access to the Internet, or if the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. For details, see Configuring Cloud Proxies in VMware Aria Operations.

The Azure VMware Solution account, with the configured private cloud, is added to the list.

### Generate a Client Secret for Azure VMware Solution

Create an Active Directory application and generate a client secret for the application in the Microsoft Azure portal. You must use the client secret when you configure a cloud account for Azure VMware Solution.

- Ensure that you are using Microsoft Azure Cloud.
- Ensure that you have a valid subscription in the Microsoft Azure portal with an Active Directory integration.

1. Log in to the Microsoft Azure portal.
2. Create an application and generate a secret for the application. For details, see Creating an Azure AD application and service principal that can access resources.

   Complete the following tasks:

   a) Create an Azure Active Directory application.

      **NOTE**
      - Ensure that the API Permission is 'Microsoft Graph User.Read'. Reader role for the Azure subscription with AVS private clouds deployed.
      - Custom role with read permissions on resource providers 'Microsoft.AVS', 'Microsoft.VMware' and 'microsoft.connectedvmwarevsphere'.

   b) Navigate to **Subscriptions** and select your subscription.
   c) In the left pane, click **Access Control (IAM)** and then click, **Add › Add Role Assignment**. Select the role you want to assign to the application. The minimum requirement is 'Reader' or above.

      **NOTE**
      To provide access to a specific resource(s), create a Resource Group for the resource(s) and give access at the resource group level.

   d) Click **Select members**, and in the right pane, search for and add one or more members you want to assign to the role for the resource.
   e) Click **Review + assign**.
   f) Generate a client secret for the application. For details, see Creating an Azure AD application and service principal that can access resources.
   g) Copy the subscription ID, directory (tenant) ID, application (client) ID, and client secret to use in your cloud account.

### Known Limitations

Review the following list of feature limitations of Azure VMware Solution integration.

- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters might appear lower than expected and capacity remaining may appear higher than expected.
- Cost drivers to calculate cost for the operating system license and additional cost per entity types are not supported.
- Cost calculation based on reference database is supported on Azure VMware Solution.
- The end-user on the vCenter on Azure VMware Solution has limited privileges. In-guest memory collection using VMware tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to VMware Aria OperationsVMware Cloud Foundation Operations using the credentials of the vCenter on Azure VMware Solution.
- The vCenter on Azure VMware Solution does not support the VMware Aria OperationsVMware Cloud Foundation Operations plugin.
- Workload optimization including pDRS and host-based business intent is not supported because the end-user does not have respective privileges to manage cluster configurations.

## Google Cloud VMware Engine

Google Cloud VMware Engine provides infrastructure as a service. It uses the scale and flexibility of the public cloud, while providing a private cloud like operating environment.

### Configuring a Google Cloud VMware Engine Instance in VMware Aria OperationsVMware Cloud Foundation Operations

You must configure a dedicated cloud account for Google Cloud VMware Engine to manage your Google Cloud VMware Engine instances in VMware Aria OperationsVMware Cloud Foundation Operations. For a successful configuration, the cloud account requires a Google Cloud Platform (GCP) project ID, service account JSON for the service account with appropriate privileges, and an optional CSP refresh token. The CSP refresh token is required if you would like to use bill-based costing and have purchased Google Cloud VMware Engine through VMware. Private clouds are auto-discovered after you save the cloud account for Google Cloud VMware Engine. You can then configure the credentials to monitor the underlying vCenter/vSAN and optionally the NSX and service discovery for each of the Private Clouds.

- Create a service account in Google Cloud Platform with at-least the viewer role privileges, note down the Google Cloud Platform project ID that you would like to manage from VMware Aria OperationsVMware Cloud Foundation Operations. Refer to the following Google Cloud Platform documentation pages for more information: Creating and Managing Service Accounts
- Generate an optional CSP refresh token for bill-based costing in the VMware Cloud Services Portal (CSP). Navigate to **API Tokens** under **My Account** and generate a CSP API token with the billing read-only role for the Google Cloud VMware Engine service.

1. From the left menu, click **Administration › Integrations**.
2. On the **Accounts** tab, click **Add**.
3. On the **Accounts Types** page, click Google Cloud VMware Engine.
4. Enter a display name and description for the cloud account.
   - Name. Enter the name for the Google Cloud VMware Engine instance as you want it to appear in VMware Aria OperationsVMware Cloud Foundation Operations.
   - Description. Enter any additional information that helps you manage your instances.
5. Enter the Google Cloud Project ID in which Google Cloud VMware Engine service has been deployed.

Google Cloud projects form the basis for creating, enabling, and using all Google Cloud services including managing APIs, enabling billing, adding and removing collaborators, and managing permissions for Google Cloud resources. Google Cloud projects are uniquely identified by an ID called Project ID. Refer to the following Google Cloud documentation for more information: Creating and Managing Projects.

6. To add credentials for the Google Cloud VMware Engine instance, click the **Add** icon, and enter the required credentials.
   - Credential Name: The name by which you are identifying the configured credentials.
   - Service Account JSON: Create a service account in Google cloud with at least the "viewer" role privileges and download its private key as a JSON file. Enter the contents of the JSON file in this field.

     **NOTE**
     You can create and use a single service account JSON that is common, similar to a super user account, for all the projects.

   - **(Optional)** CSP Refresh Token: Enter the API token if you want to use bill-based costing and Google Cloud VMware Engine was purchased from VMware. You can generate the CSP API refresh token from the Cloud Services Portal (CSP) with at least the billing read-only role for the Google Cloud VMware Engine service.

     **NOTE**
     Configure **all** the projects that are linked to the organization for accurate bill based costing.

     **NOTE**
     If any project of the Google Cloud VMware Engine adapter instance is configured without the CSP token, then reference or rate card based costing will occur.

   - Proxy Host/IP: A remote proxy server IP.
   - Proxy Port: The port that is activated on a remote proxy server.
   - Proxy Username: Enter the username of the proxy server or if you want to add a domain configured remote proxy server, then enter the username as `username@domain name`.
   - Proxy Password: Password for the proxy server username.

7. Click **Validate Connection** to validate the connection.
8. Determine which VMware Cloud Foundation OperationsVMware Cloud Foundation Operations collector or collector group is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.

   **NOTE**
   It is recommended that use cloud proxy. Ensure that there is access to the Internet and it can reach the Google Cloud VMware Engine Private Cloud's vCenter and NSX FQDNs. If the outbound internet access for the cloud proxy must be restricted, ensure the minimum cloud proxy prerequisites are met. Ensure that you have Internet connectivity for the collectors to work. For more details see, Configuring Cloud Proxies in VMware Aria Operations.

   **NOTE**
   If you have installed cloud proxy in an Google Cloud VMware Engine instance, the cloud proxy may not have outbound internet access to reach the VMware Cloud Foundation Operations service. To activate outbound internet access for the deployed cloud proxy, follow the steps described in the Google Cloud documentation in the following topic: Configuring Internet Access for Workload VMs.

9. Under Advanced Settings, enter the following details:
   - **(Optional)** Configuration Limits File Name: The Google Cloud VMware Engine account uses the following default configuration maximum file: `gcve_config_limits`. This file contains the Google Cloud VMware Engine configuration maximum soft and hard limits, and their configured value in VMware Aria

OperationsVMware Cloud Foundation Operations. If you have increased the limits for any of the Google Cloud VMware Engine configurations, you must create a new configuration file (from **Operations › Configurations › Management Packs Configuration**) and update the name of the new configuration file in this field.

- Billing Enabled: Set the option to **true** to enable bill-based costing.

10. Click **Save**.
    The page to configure the Private Clouds in Google Cloud VMware Engine appears.
11. From the list of available Private Clouds that are linked to the project configured in the Google Cloud VMware Engine instance, click any one of the Private Clouds that you want to monitor from VMware Aria OperationsVMware Cloud Foundation Operations.
12. Configure the vCenter adapter:

    1. Click the **vCenter** tab, and enter the required credentials.
       – Credential Name. The name by which you are identifying the configured credentials.
       – User Name. The vCenter user name. Use a user with the 'cloudadmin' role which has full visibility to vCenter. Users with less privileges have limited visibility, for example, the read-only users do not have visibility into management VMs.
       – Password. The vCenter password configured for that vCenter user name.
    2. Select the required collector group.

       **NOTE**
       If you have direct connectivity with your VMware Cloud vCenter, select **Default collector group**. If you are using a private IP for your vCenter or if you want to deploy Telegraf agents for application monitoring, select **Cloud Proxy**. The best practice is to deploy the Cloud Proxy on each Private Cloud instance of Google Cloud VMware Engine.

       Select the cloud proxy deployed on the given vCenter and ensure it has access to the Internet. If the outbound internet access for the cloud proxy must be restricted, ensure that the minimum cloud proxy prerequisites are met.

       For details, see Configuring Cloud Proxies in VMware Aria Operations.

       It is advised not to use the default collector groups as the Google Cloud VMware Engine management gateway firewall rule does not allow traffic originating from any address.

       If you have configured an HTTP proxy on your VMware Aria OperationsVMware Cloud Foundation Operations cloud proxy, ensure that your HTTP proxy has an exception to access the NSX Management Policy endpoint.

    3. If you have installed cloud proxy in an Google Cloud VMware Engine Private Cloud, the cloud proxy may not have outbound internet access to reach the VMware Aria OperationsVMware Cloud Foundation Operations service. To activate outbound internet access for the deployed cloud proxy, follow the steps as described in the Google documentation in the following topic: Configuring Internet Access for Workload VMs.

13. Configure the vSAN Adapter.

    1. Click the **vSAN** tab. By default, the vSAN adapter is activated.
    2. Select **Use alternate credentials** to add alternate credentials. Click the plus icon, and enter the credential name, vCenter username, and password, and click **OK**.
    3. Select **Enable SMART data collection**, if required.
    4. Click **Validate Connection** to validate the connection.

5. Click **Next**.

14. Configure the NSX adapter.

    1. Click the NSX tab and the enter the required credentials.
    2. Activate NSX configuration if it is deactivated.
    3. Click the **Add** icon next to the **Credential** field and enter the required credentials.
       - Credential Kind: Select either the NSX client certificate credential option or NSX credentials.
       - Credential Name: The name by which you are identifying the configured credentials.
       - User Name: The user name of the NSX instance if you have selected NSX credentials as the credentials kind.
       - Password: The password of the NSX instance if you have selected NSX credentials as the credentials kind.
       - Client certificate data: Enter client certificate data if you have selected NSX client certificate credentials as the credentials kind.
       - Client key data: Enter client key data if you have selected NSX client certificate credentials as the credentials kind.
       - Click **OK**.
    4. Click **Validate Connection** to validate the connection.

15. (Optional) Configure Service Discovery. For more information, see Configure Service and Application Discovery.
16. Click **Save This Private Cloud**.

    After the adapters and cloud accounts are configured, VMware Aria OperationsVMware Cloud Foundation Operations discovers and monitors the environment that runs on Google Cloud VMware Engine.

#### Known Limitations

Review the following list of feature limitations of Google Cloud VMware Engine integration.

- Management VMs are hidden from end-user visibility, hence their CPU and memory utilization are not included in the utilization of hosts, clusters, and upper level objects. As a result, the utilization of hosts and clusters may appear lower than expected and capacity remaining may appear higher than expected.
- Cost drivers to calculate cost for the operating system license and additional cost per entity types are not supported.
- Bill-based costing is supported if Google Cloud VMware Engine was purchased through VMware. Customers can use rate card costing if Google Cloud VMware Engine was purchased from a partner or a vendor other than VMware.
- The end-user on the vCenter on Google Cloud VMware Engine has limited privileges. In-guest memory collection using VMware Tools is not supported with virtual machines. Active and consumed memory utilizations continue to work in this case.
- You cannot log in to VMware Aria OperationsVMware Cloud Foundation Operations using the credentials of the vCenter on Google Cloud VMware Engine.
- The vCenter on Google Cloud VMware Engine does not support the VMware Aria OperationsVMware Cloud Foundation Operations plugin.

## VMware Cloud Foundation

In VMware Cloud Foundation, a workload domain is a policy- based resource construct with specific availability and performance attributes. It combines compute (vSphere), storage (VMware vSAN), networking (NSX), and cloud management (VMware Aria Suite) to form a single consumable entity that creates logical resource pools across compute, storage, and networking.

A workload domain consists of one or more vSphere clusters, provisioned automatically by the Domain Manager. There are two types of workload domains - the management domain and the Virtual Infrastructure (VI) workload domains.

The management domain contains the Cloud Foundation management components which include an instance of vCenter and a three-node NSX Manager cluster for the management domain. It uses the vSAN storage.

The Virtual Infrastructure (VI) workload domain is created for user workloads. For each VI workload domain, you can choose the storage option (vSAN, NFS, or VMFS on FC). A VI workload domain can consist of one or more vSphere clusters. Each cluster must start with a minimum of three hosts and can scale up to a maximum of 64 hosts. The domain manager automates the creation of the workload domain and the underlying vSphere cluster(s).

For the first VI workload domain in your environment, the SDDC Manager deploys a vCenter and a NSX Manager cluster in the management domain. For each subsequent VI workload domain, the SDDC Manager deploys an additional vCenter. New VI workload domains can share the same NSX Manager cluster with the existing VI workload domain, or deploy a new NSX Manager cluster. However, VI workload domains cannot share the management domain's NSX Manager cluster.

Configure the VMware Cloud Foundation account in VMware Aria OperationsVMware Cloud Foundation Operations to monitor these constructs of VMware Cloud Foundation.

**Configuring VMware Cloud Foundation Cloud Account in VMware Aria OperationsVMware Cloud Foundation Operations**

You must configure a VMware Cloud Foundation (VCF) cloud account to monitor a VMware Cloud Foundation environment. After configuration, all VMware Cloud Foundation domains are automatically discovered along with their connection details and credentials of underlying adapters for vCenter, vSAN, and NSX. You must save the domains want to monitor.

The VMware Cloud Foundation is not activated by default. You can activate it from the **Administration › Integrations › R epository** page or from the **Administration › Integrations › Add Accounts** page.

1. From the left menu, click **Administration › Integrations**.
2. On the Accounts tab, click **Add**.
3. On the Accounts Types page, click **VMware Cloud Foundation**.
4. Enter a display name and description for the cloud account.
   - Name. Enter the FQDN of the SDDC Manager for the VMware Cloud Foundation instance.
   - Description. Enter any additional information that helps you manage your instances.
5. Select the **Physical Data Center** you want to associate with the VMware Cloud Foundation cloud account.

   If no physical data center is created or if you want to create a new physical data center for your cloud account, you can add a new physical data center. For more information, see Adding Physical Data Centers in VMware Aria OperationsVMware Cloud Foundation Operations
6. Enter the FQDN for the **SDDC Manager** you are trying to connect.
7. To add credentials for the VMware Cloud Foundation instance, click the **Add** icon, and enter the required credentials.

   The credentials must be associated with the admin role in VMware Cloud Foundation.

   **NOTE**
   If you are using an AD or domain integrated account for authenticating your VMware Cloud Foundation instance, ensure you enter the username as username@domain or domain\username.
8. Determine which VMware Cloud Foundation Operations collector or collector group or cloud proxy is used to manage the cloud account. If you have multiple collectors or collector groups in your environment, and you want to

distribute the workload to optimize performance, select the collector or collector group or cloud proxy to manage the adapter processes for this instance.

9. Click **Validate Connection** to validate the connection.

   The Review and Accept Certificate wizard appears. Click **Accept** to validate the continue the validation.

10. Click **Advanced Settings** and enter the name of the **VCF Configuration Limit File Name**.
11. Click **Save**.
    The page to configure the Domains in VMware Cloud Foundation appears.
12. Configure a vCenter Cloud account. For more information, see Configure a vCenter Cloud Account in VMware Aria OperationsVMware Cloud Foundation Operations.

    **Important**: When you configure the vCenter, the Credential is auto-generated for VMware Cloud Foundation.

13. Configure a vSAN Adapter instance. For more information, see Configure a vSAN Adapter Instance.
14. Configure the NSX adapter. For more information, see Configuring the NSX Adapter.

    **Important**: When you configure the NSX adapter, the Credential is auto generated for VMware Cloud Foundation.

15. Configure Service Discovery. For more information, see Configure Service and Application Discovery.

    After the adapters and cloud accounts are configured, VMware Aria OperationsVMware Cloud Foundation Operations discovers and monitors the environment that runs on VMware Cloud Foundation.

16. Click **Save this SDDC**.

The VMware Cloud Foundation account, with the configured Domain, is added to the list.

You can view the operations of your VMware Cloud Foundation accounts and their domains from the VMware Cloud Foundation Operations page. For more information, see the 'Monitoring VMware Cloud Foundation (VCF) Operations' topic in the *VMware Aria OperationsVMware Cloud Foundation Operations User Guide*.

## VMware Infrastructure Health

The VMware Infrastructure Health monitors the VMware cloud management plane applications and provides metrics for their health and efficiency. With its dashboards, you can track monitor the overall health and configuration of the accounts, their applications, and services. Applications discovered by this management pack are listed in the applications section.

VMware Infrastructure Health is activated by default and a new VMware Infrastructure Health account is created for each new node and cloud proxy in VMware Aria Operations. The VMware Infrastructure Health account monitors the health of the VMware cloud management plane accounts including VMware Cloud Foundation. VMware Infrastructure Health also collects licensing information from available vCenter systems that are linked to VMware Aria Operations. For more information on licensing, see Managing Licenses

VMware Infrastructure Health in VMware Aria Operations monitors the following accounts.

1. VMware Cloud Foundation Operations for logs
2. VMware Aria Operations for Networks
3. vCenter
4. NSX
5. VMware vSAN
6. VMware Aria Automation Orchestrator
7. VMware Aria Automation
8. VMware Aria OperationsVMware Aria Operations
9. VMware Site Recovery Manager

10. VMware Identity Manager
11. VMware Cloud Foundation

The VMware Infrastructure Health focusses mainly on the operations of the VMware Cloud Foundation and is integrated with it to provide an understanding of the availability, services, virtual machines, certificates, passwords, and active alerts through the VCF Operations page. You can view the account specific and domain specific data for VMware Cloud Foundation.

**Limitations**

Starting VMware Aria OperationsVMware Cloud Foundation Operations 8.12, VMware Infrastructure Health monitors VMware Aria Operations for Networks and VMware Aria Automation accounts that are in warning state. VMware Infrastructure Health does not monitor any other accounts that are in the warning state.

VMware Infrastructure Health does not monitor accounts that are in the warning state.

### Monitoring the Health of VCF Deployments

Use VMware Infrastructure Health to monitor the health of the VMware Cloud Foundation (VCF) cloud account in VMware Aria OperationsVMware Cloud Foundation Operations. You can also monitor the VCF applications along with their services.

The domains you want to monitor should have their vCenter and NSX configured in each domain. For more information, see Configuring VMware Cloud Foundation Cloud Account. All VMware Cloud Foundation domains are automatically discovered along with their connection details and the credentials of the underlying adapters for vCenter systems and NSX.

VMware Infrastructure Health collects data related to availability, services, virtual machines, certificates, passwords, and active alerts for each configured VMware Cloud Foundation cloud account and its domains. You can view this data from the **Operations › VCF Appliances Health** page. For more information, see the 'Monitoring VMware Cloud Foundation (VCF) Appliances Health' topic in the *VMware Aria OperationsVMware Cloud Foundation Operations User Guide*.

> **NOTE**
> For VMware Cloud Foundation version 5.2 and above, NTP and DNS metrics are collected for all applications. Connectivity metrics are also collected and can be categorized as follows:
> - API connectivity is available for vCenter and NSX applications only.
> - SSH connectivity is available for all applications except NSX.
> - API and SSH Connectivity is available for host systems which are related to VCF based vCenter systems only.
>
> For VMware Cloud Foundation version 4.5.1 and above, certificates and password data will be collected for the SDDC Manager, vCenter, and NSX applications only. For versions below 4.5.1, no data will be displayed for certificates and passwords.

**Monitoring the Management Domains**

VMware Infrastructure Health will collect the following applications of the VMware Cloud Foundation management domain:

1. VMware Cloud Foundation Operations for logs application
2. vCenter application
3. VMware Identity Manager application
4. SDDC Manager application
5. NSX application
6. VMware vSAN application
7. VMware Aria Automation application

8. Fleet Management application
9. VMware Aria Operations application

**Monitoring the Workload Domains**

VMware Infrastructure Health will collect the following applications of the VMware Cloud Foundation workload domain:

1. vCenter application
2. NSX application
3. VMware vSAN application

### Monitoring vCenter Services

You can monitor the health of vCenter services with VMware Infrastructure Health.

**Resources Monitored for vCenter**

VMware Infrastructure Health starts collecting the health of the following vCenter services:

- vCenter Appliance Health Services
- vCenter Health Services
- vCenter Backup job
- vCenter NTP server
- vCenter Licensing Group

> **NOTE**
> To collect the vCenter health services on vCenter version 8.0.3 and above, you must login to vCenter VAMI client and then start the VMware Service Lifecycle Manager API. The start type must be set to automatic in vCenter Server Management version 8.0.3 and above.
>
> The list of available services differ based on the type of Operating System and vCenter version.
>
> The NTP and Backup Jobs services are available only for VCVA 6.5 version and above.

**Resources Monitored for VMware vSAN**

VMware Infrastructure Health starts collecting the health of the following VMware vSAN services:

- VMware vSAN Health services

### Permissions Required to Discover vCenter Services

A user must have certain privileges to discover the services of vCenter.

The user should be a member of the 'SystemConfiguration.Administrator' group or have the administrator permissions to discover the vCenter services.

### vCenter Certificate Monitoring

Once the VMware Infrastructure Health collects the health of the vCenter health services, it monitors the ESXi host and vCenter root certificates, on a daily basis. It also alerts the user to avoid environment downtime. The certificate properties are published under the Certificate Summary group.

### vCenter Backup Job

The retention period for vCenter backup job monitoring is set by default to the last 7 days. You can configure the retention period to get more number of backup job items.

#### Backup Job Retention

To set the backup job retention period for vCenter:

1. From the left menu, click **Administration > Integrations**, and then click **Accounts** tab.
2. On the Account tab, expand the **VMware Infrastructure Health** and click the **Vertical Ellipses > Edit**.
3. Under **Advanced Settings**, in the **vCenter Backup Job Retention Period (in Days)** field, set the period (in days) for which vCenter backup jobs should be retained.

### Monitoring NSX Services

You can monitor the health of NSX services with VMware Infrastructure Health.

#### Resources Monitored for NSX

VMware Infrastructure Health monitors the health of the following NSX services:

- http
- nsx-ui
- ntp
- syslog
- mgmt-plane-bus
- install-upgrade
- snmp
- mgmt-plane-bus
- nsx-upgrade-agent
- migration-coordinator
- syslog
- nsx-ui
- node-mgmt
- nsx-message-bus
- migration-coordinator
- ssh
- snmp
- cm-inventory
- manager
- ssh
- liagent
- ntp
- install-upgrade
- http
- controller
- cluster-boot-manager
- search
- cm-inventory
- node-mgmt

- nsx-message-bus
- nsx-upgrade-agent
- telemetry
- telemetry
- manager
- controller
- cluster-boot-manager
- liagent
- search

### NSX Certificate Monitoring

Once the VMware Infrastructure Health collects the health of the NSX health services, it monitors the configured NSX Server and NSX manager node certificates, on a daily basis. It also alerts the user to avoid environment downtime. The certificate properties are published under the Certificate Summary group.

Alerts are raised based on the number of days set for critical, immediate, and warning category under the global setting. The NSX certificates contain the following three symptoms:

- has expired or will expire shortly - for critical
- expire shortly - for immediate
- about to expire - for warning

These symptoms are triggered if the certificate expiry date has reached the threshold condition set for critical, immediate, and warning. A notification is displayed explaining the alert information.

### NSX Backup Job

The retention period for NSX backup job monitoring is set by default to the last 7 days. You cannot update the retention period for NSX.

## Monitoring Workspace ONE Access Services

You can monitor the health of VMware Workspace ONE Access services with VMware Infrastructure Health.

### Resources Monitored for Workspace ONE Access

VMware Infrastructure Health collects the health of the following Workspace ONE Access services:

- Identity Manager Cluster Node
- Cert Proxy Certificate
- Certificates
- Cert Proxy
- Disk Space
- AirWatch API Server
- Connector
- Configurator
- ACS Health
- Identity Manager FQDN
- Port Connectivity
- Database Connection
- Application Manager
- RabbitMQ
- Elasticsearch Health
- EhCache Cluster Diagnostics

- Messaging Connection
- Analytics Connection

### Monitoring VMware Aria Operations for Networks Services

You can monitor the VMware Aria Operations for Networks services with the VMware Infrastructure Health.

**Resources Monitored from VMware Aria Operations for Networks**

The Networks Platform Node and the Networks Collector Node contains services which are called Network services.

VMware Infrastructure Health monitors the health of the following VMware Aria Operations for Networks services:

- FoundationDB
- Nfcapd
- IpfixProcessor
- NetopaCollector
- CollectorMain
- All
- ntpsec
- Nginx
- SwitchTelemetryS

### Monitoring VMware Aria Automation Services

You can monitor the health of VMware Aria Automation services with the VMware Infrastructure Health.

**Resources Monitored for VMware Aria Automation**

VMware Infrastructure Health collects the health of the following VMware Aria Automation services:

- ingress-ctl
- kube-dns
- etcd-service
- health-reporting-service
- kube-apiserver
- kube-controller-manager
- kube-flannel-ds
- kube-proxy
- kube-scheduler
- kubelet-rubber-stamp
- predictable-pod-scheduler
- tiller-deploy
- openfaas
- abx-service
- approval-service
- assessment-service
- ui
- catalog-service
- cgs-service
- cmx-service
- codestream
- docker-registry
- ebs

- form-service
- hcmp-service
- identity-service
- migration-service
- no-license
- postgres
- project-service
- provisioning-service
- proxy-service
- rabbitmq-ha
- relocation-service
- tango-blueprint-service
- tango-vro
- terraform-service
- user-profile-service
- vco
- adapter-host-service
- endpoints
- lemans-resources
- lemans-gateway
- private-cloud-gateway

**NOTE**
By default, VMware Infrastructure Health monitors all VMware Aria Automation 8.x services. If you do not want certain VMware Aria Automation services to be monitored, you must add the Automation services names (such as endpoints, kube-dns, ebs, etc.) in the VMware Infrastructure Health properties file under the key **CAS_SERVICE_TO_IGNORE**.

If the VMware Aria Automation is attached to the load balancer, activate port forwarding to monitor Automation services with VMware Infrastructure Health. Refer to Knowledge Base for more details. Ignore these steps if the VMware Aria Automation is a standalone node deployment.

### Monitoring SDDC Manager Services

You can monitor the health of SDDC Manager services with VMware Infrastructure Health.

#### Resources Monitored for SDDC Manager

VMware Infrastructure Health starts collecting the health of the following SDDC Manager services:

- COMMON_SERVICES
- DOMAIN_MANAGER
- LCM
- OPERATIONS_MANAGER
- SDDC_MANAGER_UI

#### SDDC Manager Backup Job

The retention period for SDDC backup job monitoring is set by default to the last 7 days. You can configure the retention period to get more number of backup job items.

Configure the backup job retention period for SDDC Manager in VMware Infrastructure Health.

1. From the left menu, click **Administration > Integrations**, and then click the **Accounts** tab.

2. On the Account tab, expand the **VMware Infrastructure Health** and click the **Vertical Ellipses > Edit**.
3. Under **Advanced Settings**, in the **SDDC Manager Backup Job Objects Retention Period (in Days)** field, set the period (in days) for which SDDC Manager backup jobs should be retained.

## OS and Application Monitoring

You can monitor application services in VMware Aria OperationsVMware Cloud Foundation Operations. You can use the product-managed agent to monitor physical servers and VMs that are managed or unmanaged by the vCenter cloud account (target machine).

For example, as an administrator, you might need to ensure that the infrastructure provided for running the application services is sufficient and that there are no problems. If you receive a complaint that a particular application service is not working properly or is slow, you can troubleshoot by looking at the infrastructure on which the application is deployed. You can view important metrics related to the applications and share the information with the team managing the applications. You can use VMware Aria OperationsVMware Cloud Foundation Operations to deploy the agents and send the related application data to VMware Aria OperationsVMware Cloud Foundation Operations. You can view the data in VMware Aria OperationsVMware Cloud Foundation Operations and share it with the team so that they can troubleshoot the application service.

You can conduct remote checks, monitor Windows services and Linux processes, monitor operating systems and applications, and run custom scripts in VMware Cloud Foundation Operations.

Using VMware Aria Operations Advanced edition, you can monitor operating systems, conduct remote checks, custom script, and monitor Windows services and Linux processes in VMware Aria Operations. Using VMware Aria Operations Enterprise edition gives you additional capabilities to monitor operating systems and applications, and run custom scripts in VMware Aria Operations. Monitoring AWS and Azure instances through the Telegraf agent is not supported in any of the on-prem license editions.

### Product-Managed Telegraf on Different Types of Machines

The following (object types) are now supported for application monitoring based on the type of machine the product-managed Telegraf agent is running on.

- `Virtual Machine`: Refers to the vCenter. All the capabilities of product-managed Telegraf agents and monitoring of the vCenter VMs remain the same from previous releases. For more information, see Install and Uninstall a Telegraf Agent and Additional Operations from the Manage Telegraf Agents Page .
- `Endpoint`: VMware Aria OperationsVMware Cloud Foundation Operations additionally supports the following types of machines that are categorized under an object type called `Endpoint`:
  - Physical servers
  - Machines that are not directly controlled by VMware Aria OperationsVMware Cloud Foundation Operations. For example, VMs that are not managed by the vCenter cloud account.
    To install and uninstall a product-managed Telegraf agent on a remote machine or a physical server, use the helper script. For more information, see Install/Uninstall an Agent Using a Script on a Linux Platform and Install/Uninstall an Agent Using a Script on a Windows PlatformInstall/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform . To manage the life cycle of agents on a remote machine or a physical server, use the VMware Aria OperationsVMware Cloud Foundation Operations UI or suite-api. For more information, see Additional Operations from the Manage Telegraf Agents Page .

  In the application monitoring documentation for product-managed Telegraf agents, the term **target machine/s** refers to physical servers and the VMs that are managed and unmanaged by the vCenter cloud account.

### High Availability for Application Monitoring

You can now use application monitoring high availability activated collector groups in failover mode to monitor applications without loosing any data (metrics) from the application that is monitored even if your cloud proxy is down or if any of the cloud proxy components fail. Here are some important points:

- Cloud proxies in an application monitoring high availability activated collector group work in a primary-secondary mode.
- You cannot convert an existing Telegraf agent on all target machines, to high availability. To convert an existing target machine to high availability, you must install or re-install the agent with an application monitoring high availability activated collector group.
- You will lose all application data when you add a cloud proxy used by Telegraf agents, to the application monitoring high availability activated collector group.
- Failover or failback has a downtime of three collection cycles.
- Addition or removal of a cloud proxy from the application monitoring high availability activated collector has a downtime of three collection cycles.

**Steps to Activate Application Monitoring High Availability on Collector Groups**

Here are the steps to activate application monitoring high availability on collector groups and use it for application monitoring and data collection:

1. Activate high availability for specific cloud proxies when you add a collector group. For more information, see Adding a Collector Group.
2. (Optional) You can view the cloud proxies that are configured for high availability from the **Cloud Proxies** page. For more information, see Monitoring the Health of Cloud Proxies.
3. While installing an agent, select a collector group that has application monitoring high availability activated. For more information, see Install an Agent From the UI.
4. (Optional) You can view the following agent details such as:

    a. If high availability is activated for the collector group (Collector Group column).
    b. The virtual IP configured on an application monitoring high availability activated collector group to which all target machines send data (Virtual IP details column).
    c. The IP address of the primary cloud proxy that is collecting data (Cloud Proxy column).

    You can view these details from the **Manage Telegraf Agents** page (**Operations › Applications › Manage Telegraf Agents**). For more information, see Additional Operations from the Manage Telegraf Agents Page.

**Supported Application Services**

OS and Application monitoring in helps virtual infrastructure administrators and application administrators to discover operating systems and applications running in provisioned guest operating systems at a scale and to collect run-time metrics of the operating system and application for monitoring and troubleshooting respective entities.

The following 23 application services are supported.

| Application Services |
|---|
| Active Directory |
| Active MQ |
| Apache HTTPD |
| Cassandra |
| HyperV |
| Java Application |
| JBoss Server |
| MongoDB |
| MS Exchange |
| Microsoft IIS |

*Table continued on next page*

*Continued from previous page*

| Application Services |
|---|
| Microsoft SQL Server |
| MySQL |
| Network Time Protocol |
| Nginx |
| Oracle DB |
| Pivotal TC Server |
| PostgreSQL |
| RabbitMQ |
| Riak KV |
| SharePoint Server |
| Tomcat Server |
| Oracle WebLogic Server |
| WebSphere |

**VeloCloud Application Services**

Eight VeloCloud application services are supported in VMware Aria OperationsVMware Cloud Foundation Operations. The supported application services are listed here. Some of the application services have mandatory properties which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected. The following services are supported in VMware Aria OperationsVMware Cloud Foundation Operations:

| Application Services |
|---|
| VeloCloud Orchestrator |
| Nginx |
| Clickhouse |
| Network Time Protocol |
| MySQL |
| Redis |
| Java Application |
| VeloCloud Gateway |

### Supported Platforms

VMware Aria OperationsVMware Cloud Foundation Operations supports monitoring for the following platforms and app combinations with API support.

**Platforms Supported by VMware Aria OperationsVMware Cloud Foundation Operations for OS and Application Monitoring**

| Platform | Version | Architecture | Application |
|---|---|---|---|
| Red Hat Enterprise Linux | 7.x | 64-bit | OS Metrics and all supported applications. |

*Table continued on next page*

*Continued from previous page*

| Platform | Version | Architecture | Application |
|---|---|---|---|
| | 8.x<br>9.x | | |
| CentOS | 7.x<br>8.x | 64-bit | OS Metrics and all supported applications. |
| Oracle Linux | 7.x<br>8.x<br>9.x | 64-bit | OS Metrics and all supported applications. |
| Rocky Linux | 8.x<br>9.x | 64-bit | OS Metrics and all supported applications. |
| Ubuntu | 18.04 LTS<br>16.04 LTS<br>20.x | 64-bit | OS Metrics and all supported applications. |
| VMware Photon Linux | 1.0<br>2.0<br>3.0 | 64-bit | Only OS metrics and custom monitoring actions are supported for photon OS.<br>VMware Aria Application Remote Collector 8.3, 8.2, 8.1, and 7.5 runs on Photon 1.0.<br>Site Recovery Manager 8.2 runs on Photon 2.0<br>vSphere- vSphere 6.7 & 6.5 runs on Photon OS 1.0<br>VMware vSAN 6.7 & VMware vSAN 6.5 runs on Photon OS 1.0<br>Unified Access Gateway 3.7 runs on Photon 3.0 & 3.6 runs on Photon 2.0. |
| SUSE Linux Enterprise Server | 12.x<br>15.x | 64-bit | OS Metrics and all supported applications. |
| Windows | Windows Server 2019<br>Windows Server 2016<br>Windows 2012 | 64-bit | OS Metrics and all supported applications. |

*Table continued on next page*

*Continued from previous page*

| Platform | Version | Architecture | Application |
|---|---|---|---|
| | Windows Server 2012 R2<br>Windows Server 2022 | | |

**Prerequisites**

To monitor your application services and operating systems, complete all the prerequisites so that cloud proxy can communicate successfully with vCenter and the end points.

**Port Information and Communication with the Target Machines**

Port information and communication details with the target machines to install the Telegraf agent from the UI or via the helper script is detailed on this page.

Ensure that cloud proxy's ports are reachable from target machines both via the IP and FQDN.

For the latest port information, see VMware Ports and Protocols.

**Figure 7: Port Information and Communication with the vCenter, VMs, VMware Aria OperationsVMware Cloud Foundation Operations, and Cloud Proxy (Agent Install from the UI)**

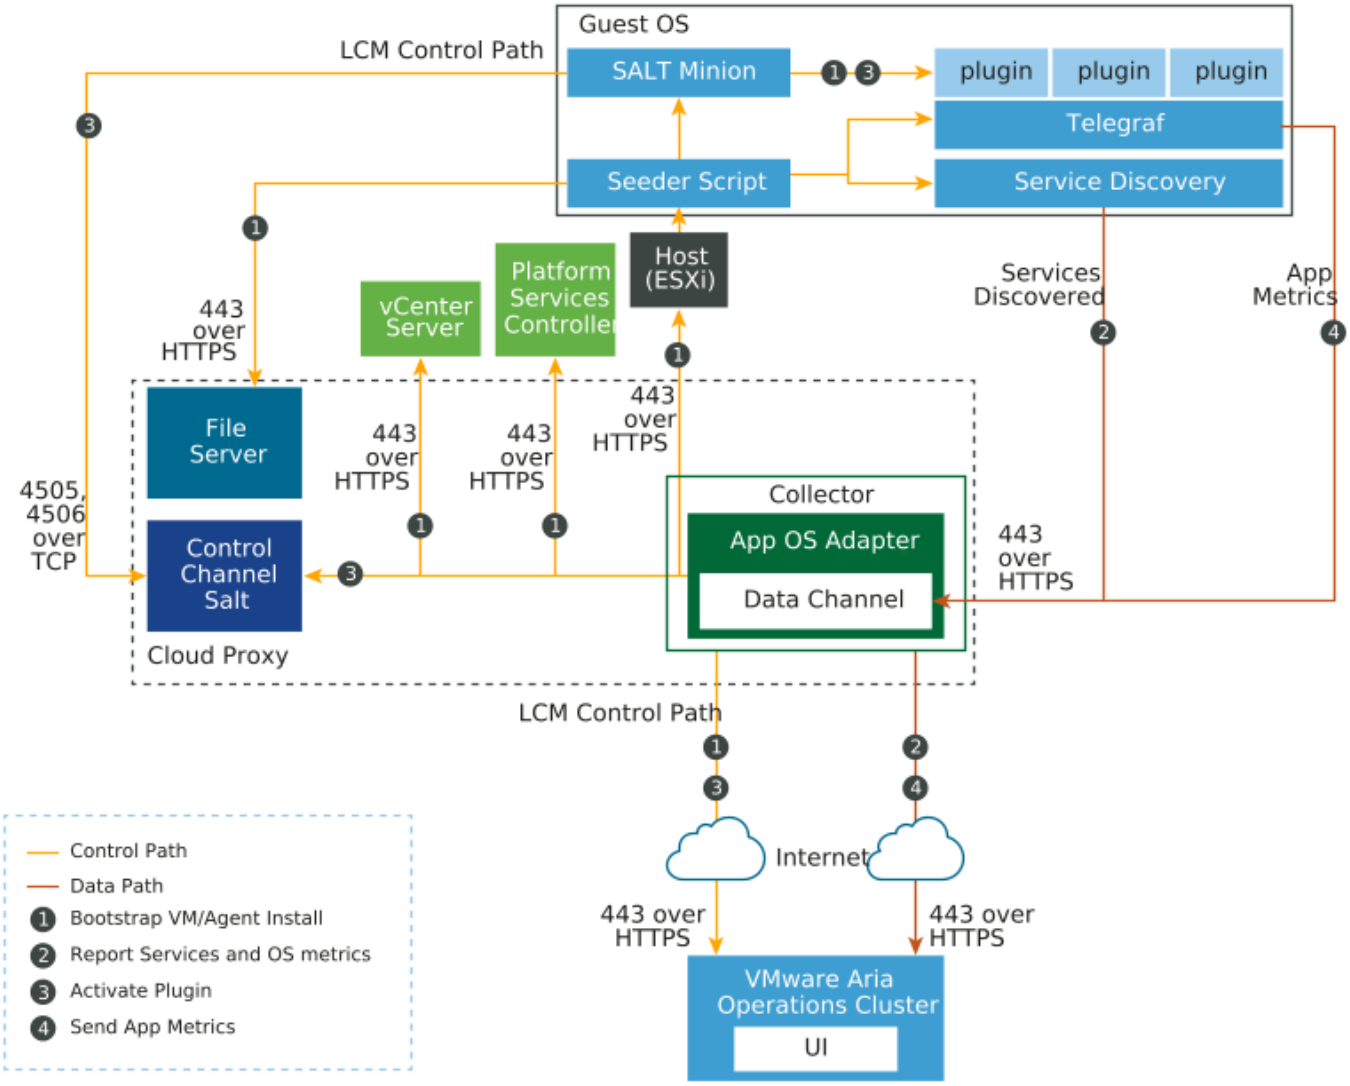

**Figure 8: Port Information and Communication with the Target Machines and VMs for Script-Based Agent Install**

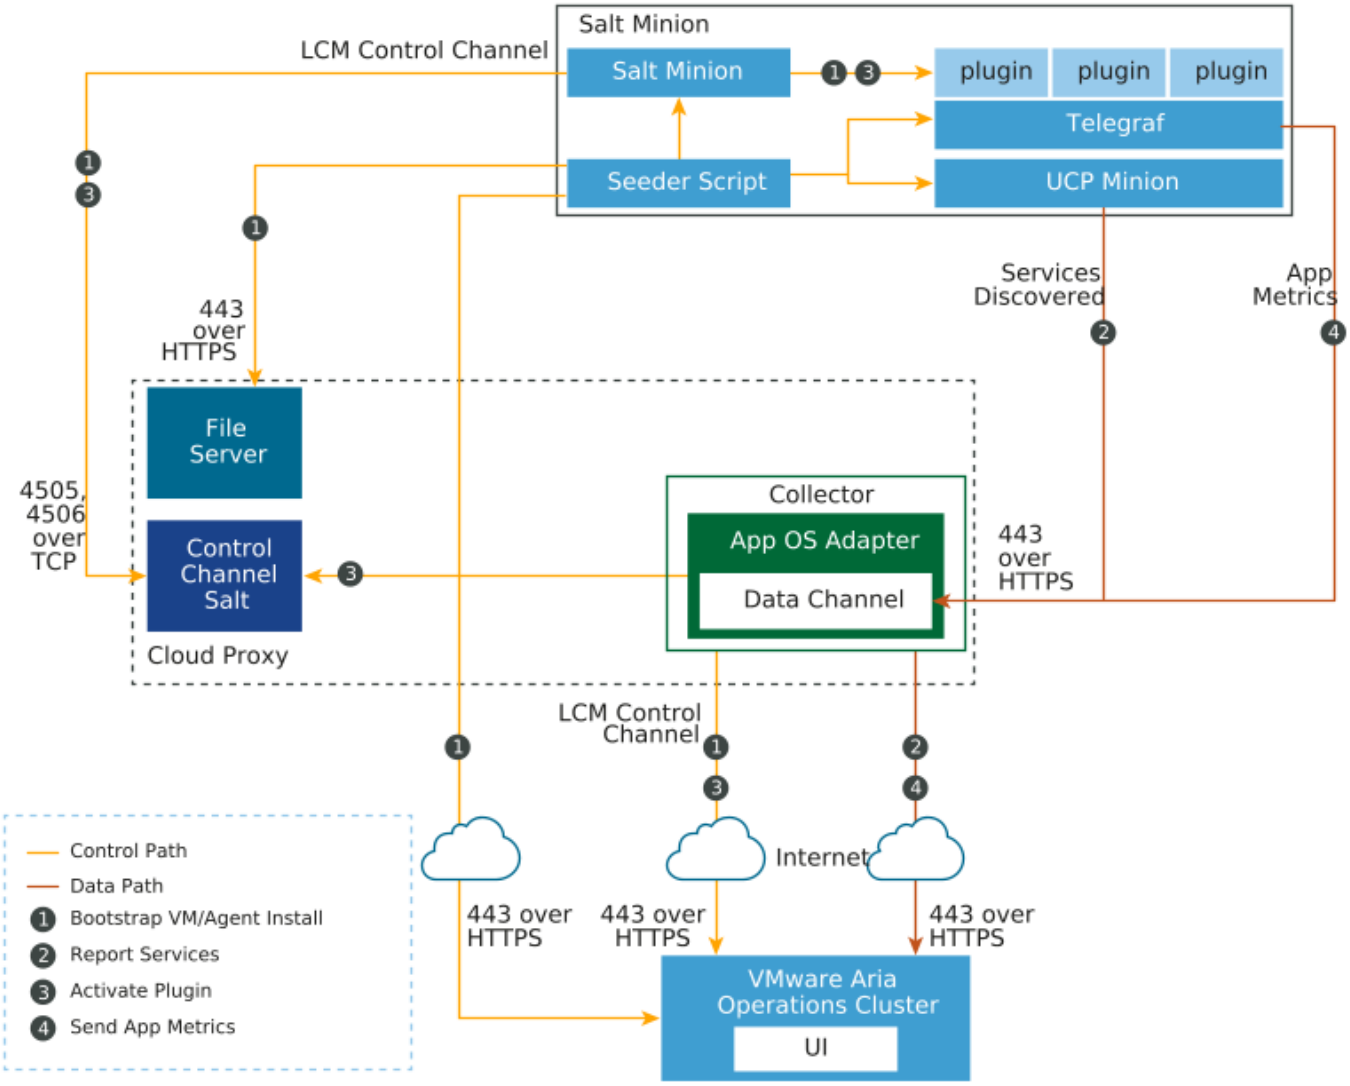

**Figure 9: Port Information and Communication with the vCenter, VMs, VMware Aria OperationsVMware Cloud Foundation Operations, and Cloud Proxy (Agent Install from the UI)**

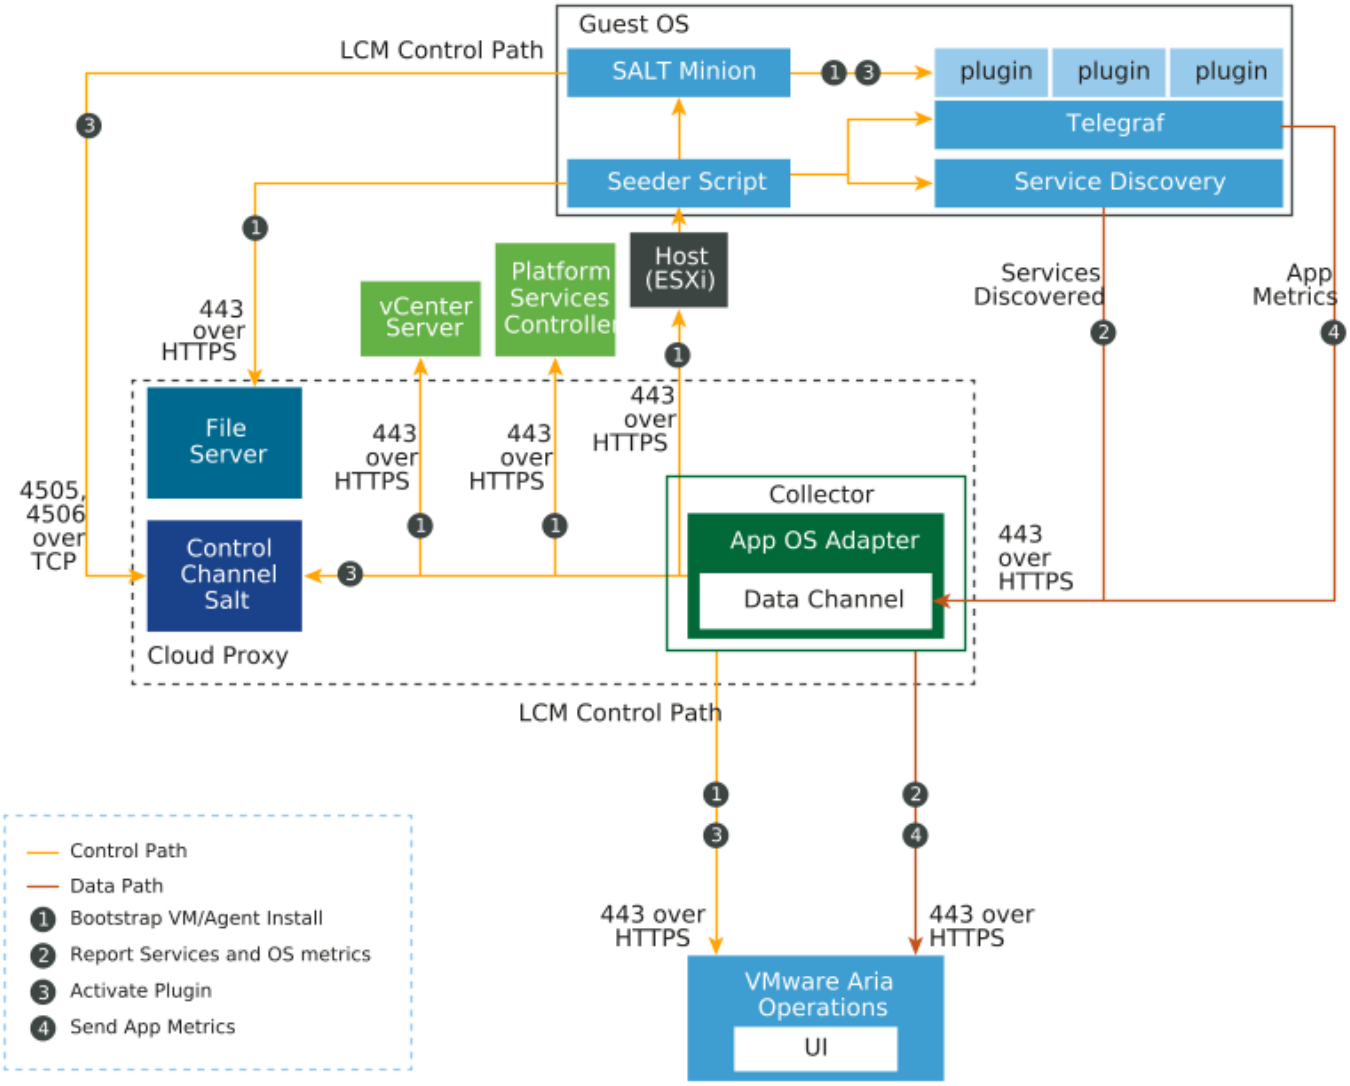

**Figure 10: Port Information and Communication with the Target Machines and VMs for Script-Based Agent Install**

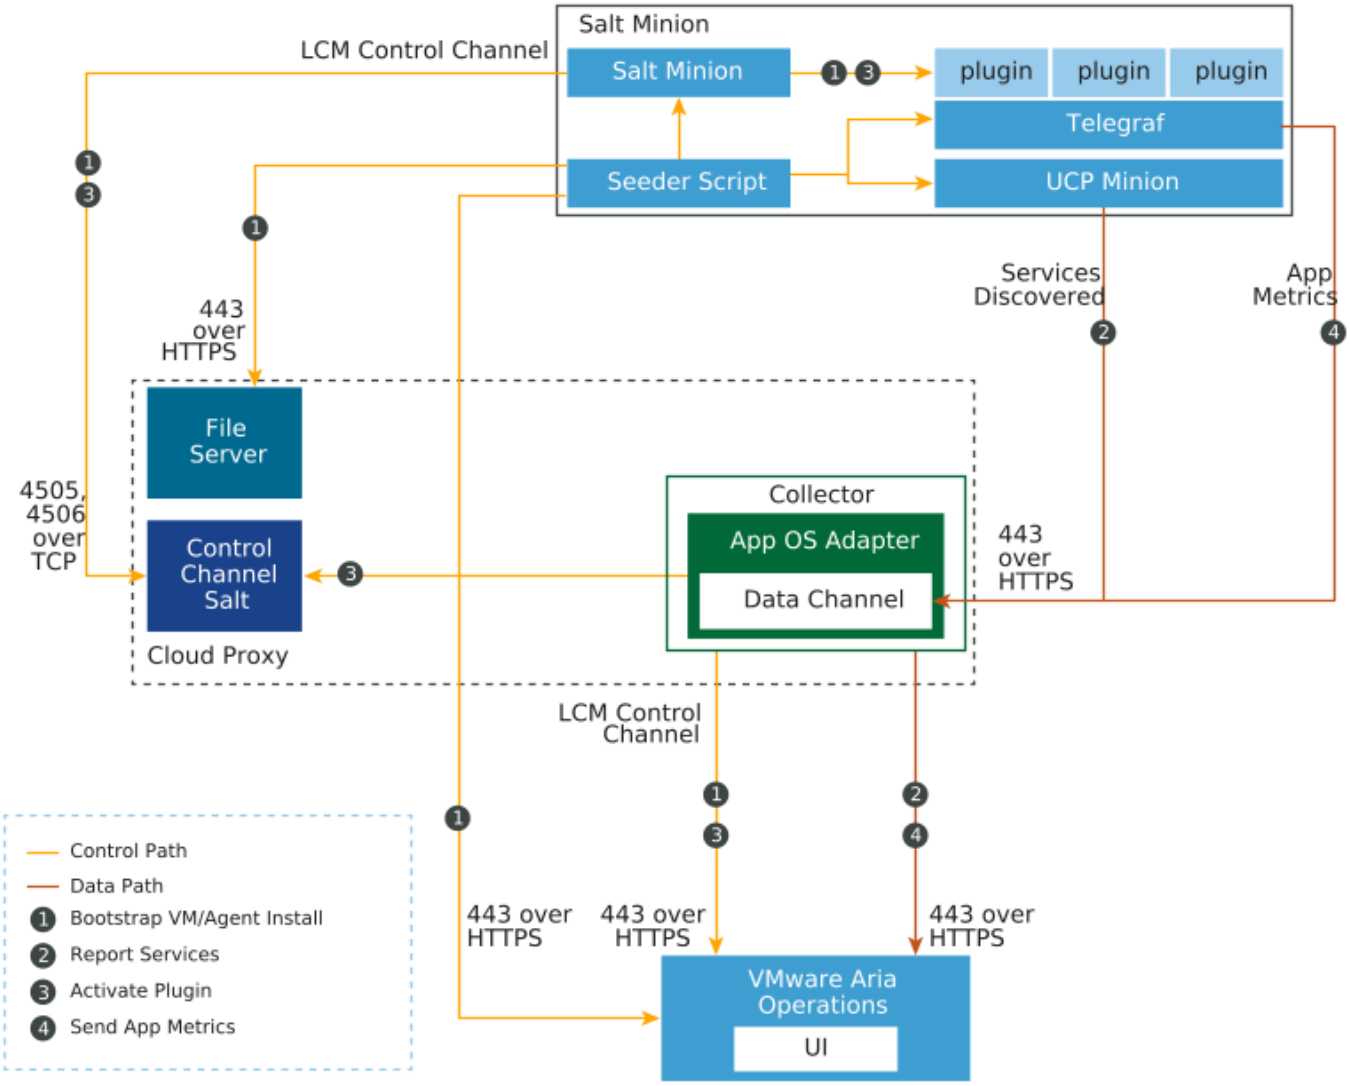


## Communication with Cloud Proxy and vCenter

Complete all the prerequisites required so that cloud proxy can communicate with vCenter. In the case of application monitoring high availability activated collector groups, these prerequisites are applicable for cloud proxies that are a part of the collector group.

> **NOTE**
> Applicable to VMs managed by the vCenter cloud account.

- Port 443 in the ESXi where the workload VMs are deployed must be accessible to cloud proxy.
- Port 443 in Platform Services Controller and vCenter is accessible to cloud proxy. Open this port if vCenter is configured with an external Platform Services Controller.
- Configure a vCenter adapter. The vCenter user should have read access at the vCenter Server level and should also have the following permissions: `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries`.

**NOTE**
For script-based agent install and uninstall, the `Guest operation modifications`, `Guest operation program execution`, and `Guest operation queries` permissions are not required.

### Communication with Cloud Proxy and Target Machines

Complete the prerequisites required during the handshake of cloud proxy with the target machines.

For the handshake of cloud proxy with the target machines, the following prerequisites are required:

- The target machine where Telegraf should run, cloud proxy, and VMware Aria OperationsVMware Cloud Foundation Operations should be time synchronized. In case of VMs managed by vCenter cloud account, the ESXi instance where the VM is deployed, vCenter, VMs, VMware Aria OperationsVMware Cloud Foundation Operations and cloud proxy should be time synchronized.
- The target machines must have access to ports 443, 4505, and 4506 on cloud proxy and on the virtual IP of the application monitoring high availability activated collector group. You can verify access by running the following commands on the target machine:
  - For Linux Machines:

```
timeout 10 bash -c "</dev/tcp/{cloudproxy_fqdn_or_virtual_IP}/4505"

echo $?

timeout 10 bash -c "</dev/tcp/{cloudproxy_fqdn_or_virtual_IP}/4506"

echo $?

timeout 10 bash -c "</dev/tcp/{cloudproxy_fqdn_or_virtual_IP}/443"

echo $?
```

  - For Windows Machines:

```
wget.exe --spider -t 1 -T 10 {cloudproxy_fqdn_or_virtual_IP}:4505

wget.exe --spider -t 1 -T 10 {cloudproxy_fqdn_or_virtual_IP}:4506

wget.exe --spider -t 1 -T 10 {cloudproxy_fqdn_or_virtual_IP}:443
```

    **NOTE**
    If you do not have `wget.exe` on the Windows machine, navigate to the `%temp%` folder or it's parent folder in File Explorer and search for `wget.exe` after you attempt installation.

    **NOTE**
    In the above commands, use virtual IP in the case of application monitoring high availability activated collector groups. For individual cloud proxy or cloud proxy which belongs to application monitoring high availability deactivated collector groups, use the cloud proxy FQDN.

- The necessary privileges for a user which are required for agent installation are mentioned in the User Account Prerequisites page.
- Target machine configuration requirements.
  - Linux requirements
    Commands: **`/bin/bash, sudo, tar, awk, curl`**

    Packages: `coreutils (chmod, chown, cat), shadow-utils (useradd, groupadd, userdel, groupdel), net-tools`

Configure mount point on `/tmp` directory to allow script execution.

  - Windows requirement
    - The Visual C++ version must be higher than 14.
    - Performance Monitors on a Windows OS VM must be activated.
  - Windows 2012 R2 requirement
    The target machine must be updated with the Universal C Runtime. Refer to the following link for more information.
- VMware Tools must be installed and running on the vCenter VM on which you want to install the agent. For information about supported VMware Tools versions, click this link.

  **NOTE**
  Applicable only to VMs managed by the vCenter cloud accounts.
- Add a `tmp` folder with "exec" permission to install agents on the latest UAG Photon OS VM. To configure a mount point on the `/tmp` directory to allow script execution, run the following command: `<mount -o remount,exec /tmp>`

### User Account Prerequisites

There are certain user account prerequisites required for the install of agents.

#### Windows Target Machine User Account Requirements

- To install agents,
  - The user must be either an administrator, or
  - A non-administrator who belongs to the administrator group.

#### Linux Target Machine User Account Requirements

For Linux target machines, there are two user accounts for the Telegraf agent, such as the install user and the run-time user. User credentials which are provided during agent installation, are for the install user. The *arcuser* is a run-time user and needs a set of privileges which are necessary for the agent's components to run.

- `/tmp` mount point should be mounted with exec mount option.
- The following are minimal necessary permissions of the user to install agents and should be mentioned in `sudoers` file:
  For example, for a user called **telegrafinstall**, you can find the `sudoers` file in the `/etc/sudoers` file or in the folder `/etc/sudoers.d/`:

```
Defaults:telegrafinstall !requiretty


Cmnd_Alias ARC_INSTALL_USER_COMMANDS=/usr/bin/cp*,/bin/cp*,/usr/bin/mkdir*,/bin/
mkdir*,/usr/bin/chmod*,/bin/chmod*,/opt/vmware/ucp/bootstrap/uaf-bootstrap.sh,/opt/
vmware/ucp/ucp-minion/bin/ucp-minion.sh


telegrafinstall ALL=(ALL)NOPASSWD: ARC_INSTALL_USER_COMMANDS
```

### Run-Time User Prerequisites

There are two ways in which a run-time user is created in Linux target machines: automatically and manually. A run-time user has a standard name and group, which is the *arcuser* and *arcgroup* respectively. If the **Create run time user on linux virtual machines, with required permissions as part of agent installation** check box is selected, the *arcuser* and *arcgroup* are created automatically. The check box is selected by default. A run-time user is also created

automatically during a script-based install. If you choose to manually create the *arcuser* and *arcgroup*, here are the steps to do it manually:

Create the *arcgroup* and *arcuser* and associate the *arcgroup* as the primary group of the *arcuser*.

1. The *arcgroup* must be the primary group of the *arcuser*.
   The following commands can be used to create the *arcgroup* and *arcuser*:

   **groupadd arcgroup**

   **useradd arcuser -g arcgroup -M -s /bin/false**
2. The *arcuser* must be created with no home directory and no access to the login shell.
   For example, the **/etc/passwd** entry for the *arcuser* is as follows after adding *arcuser* and *arcgroup*.

   **arcuser:x:1001:1001::/home/arcuser:/bin/false**
3. The **arcuser** must have password-less specific set of privileges as mentioned below, which must be written in **/etc/sudoers** file or in the folder **/etc/sudoers.d/**:

```
Defaults:arcuser !requiretty

Cmnd_Alias VAPCOMMANDS=/usr/bin/systemctl * ucp-telegraf.service, !/usr/bin/systemctl
* * ucp-telegraf.service, /bin/systemctl * ucp-telegraf.service, !/bin/systemctl * *
ucp-telegraf.service, /usr/bin/systemctl * ucp-minion.service, !/usr/bin/systemctl *
* ucp-minion.service, /bin/systemctl * ucp-minion.service, !/bin/systemctl * * ucp-
minion.service, /usr/bin/systemctl * salt-minion.service, !/usr/bin/systemctl * *
salt-minion.service, /bin/systemctl * salt-minion.service, !/bin/systemctl * * salt-
minion.service, /usr/bin/systemctl * ucp-salt-minion.service, !/usr/bin/systemctl * *
ucp-salt-minion.service, /bin/systemctl * ucp-salt-minion.service, !/bin/systemctl *
* * ucp-salt-minion.service, /usr/bin/netstat, /bin/netstat, /opt/vmware/ucp/tmp/
telegraf_post_install_linux.sh, /opt/vmware/ucp/bootstrap/uaf-bootstrap.sh, /opt/
vmware/ucp/content/runscript.sh, /opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh, /usr/
bin/systemd-run, /bin/systemd-run

arcuser ALL=(ALL) NOPASSWD: VAPCOMMANDS
```

### Steps to Monitor Applications

You can monitor and collect metrics for your application services and operating systems.

The following flowchart describes how you can set up VMware Aria OperationsVMware Cloud Foundation Operations for application monitoring in the case of VMs managed by the vCenter cloud account.

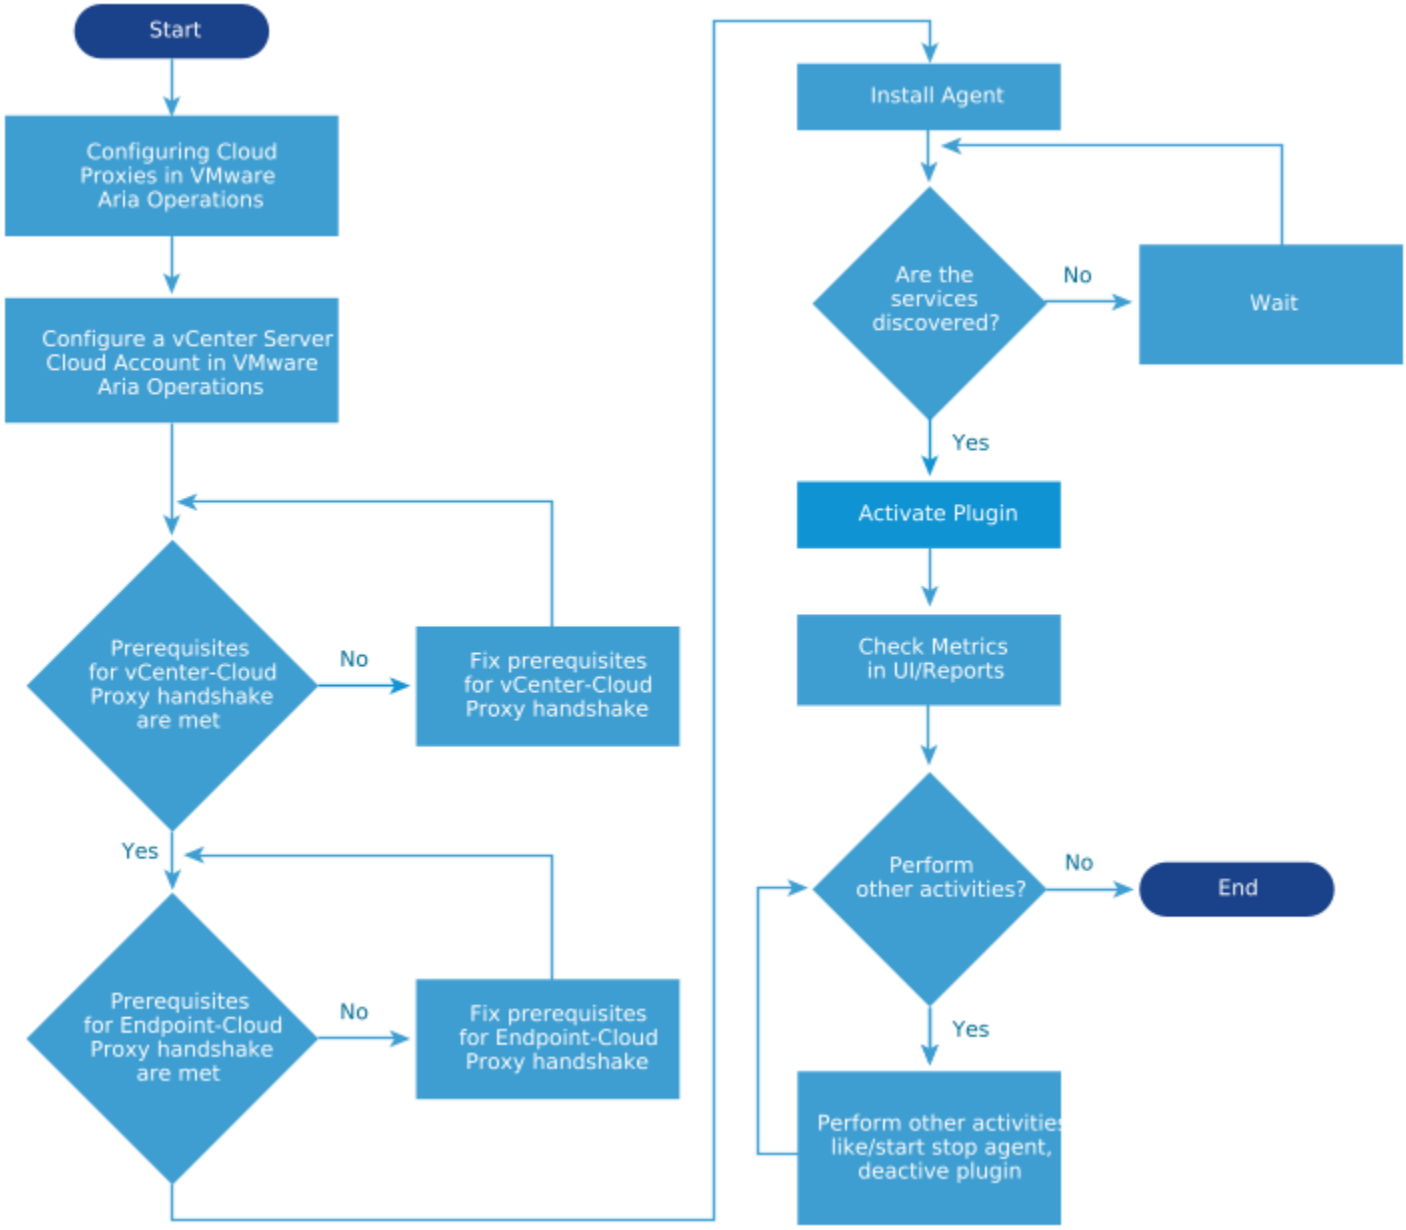


Follow these steps to monitor applications.

1. Configure cloud proxy on which the AppOS adapter instance is created.
   For more information, see Configuring Cloud Proxies in VMware Aria Operations.

   For more information, see Configuring Cloud Proxies in VMware Aria Operations.
2. Configure a vCenter cloud account that uses the cloud proxy configured in step 1.Configure a vCenter cloud account.
   For more information, see Configure a vCenter Server Cloud Account in VMware Aria Operations.

   **NOTE**
   This step is not applicable for target machines.
3. Complete all the prerequisites.
   For more information, see Prerequisites.

4. Install agents on selected target machines.
   For more information, see Install an Agent from the UI or Install/Uninstall an Agent Using a Script on a Linux Platform and Install/Uninstall an Agent Using a Script on a Windows PlatformInstall/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform .
5. Activate an application service.
   For more information, see Activate an Application Service.
6. View the summary of application services and operating systems discovered in VMware Aria OperationsVMware Cloud Foundation Operations.
   For more information about monitoring your applications in VMware Aria OperationsVMware Cloud Foundation Operations, see Summary of Discovered and Supported Operating Systems and Application Services.

### Install and Uninstall a Telegraf Agent

You can install and uninstall Telegraf agents on VMs managed by the vCenter cloud account from the user interface of VMware Aria OperationsVMware Cloud Foundation Operations or by running a script on physical servers and on machines not managed by VMware Aria OperationsVMware Cloud Foundation Operations.

#### Install an Agent from the UI

You must select the vCenter VMs on which you want to install the agent. All the vCenter VMs of the vCenter cloud account are listed in the **Manage Telegraf Agents** page.

> **NOTE**
> You cannot install an agent from the VMware Aria OperationsVMware Cloud Foundation Operations user interface, on physical servers and target machines (Discovered in VMware Aria OperationsVMware Cloud Foundation Operations as the `Endpoint` object type).

Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

1. From the left menu, select **Operations › Applications**. From the **Applications** panel, select **Manage Telegraf Agents**. You see the **Manage Telegraf Agents** page.
2. From the **Manage Telegraf Agents** page, select the VMs managed by the vCenter cloud account on which you want to install the agent.
3. From the **Manage Telegraf Agents** page, click **Agent Actions**, and then click **Install**.
4. If the VMware Aria Operations Application Monitoring Adapter instance has not been created for the specific vCenter, you see the **Installing Telegraf Agent** dialog box. You can configure high availability for application monitoring via collector groups:
   - Select **Cloud Proxy Groups, with High-Availability** if you want to associate an application monitoring high availability activated collector group to the vCenter. You can choose the collector group from the **HA Enabled Group** drop down list, or
   - Select **Cloud Proxies, without High-Availability** if you want to associate a cloud proxy to the vCenter. You can choose a cloud proxy from the **Cloud Proxy** drop down list.
   - Click **Done**.

   > **NOTE**
   > If the vCenter is on the default collector, the cloud proxy options are blank. If the vCenter is on a specific cloud proxy, the cloud proxy option is automatically populated. To deploy the VMware Aria Operations Application Monitoring Adapter instance on a different cloud proxy, select another cloud proxy or application monitoring high availability activated collector group.

5. You see the **Manage Agent** dialog box.
6. From the **How do you want to provide VM Credentials** page, complete the following steps:

a) If you have a common user name and password for all the VMs, select the **Common username and password** option.
b) If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.
c) Under the section called **Telegraf Configurations**, select the **Telegraf agent configurations - Use backup of configurations if available** option if you want to restore the configuration details that were backed up for available agents on the VM, during a previous uninstall of the agents. For more information see, Uninstall an Agent.

> **NOTE**
> Some agents such as Active MQ, Java Application, JBoss Server, Pivotal TC Server, Tomcat Server, Oracle WebLogic Server, and Riak KV might not work after you upgrade VMware Aria Operations from a previous version to 8.18 and choose to restore the agent configurations. As a workaround, after you restore agent configurations during the install, select the agents that were restored after you upgraded VMware Aria Operations to 8.18 and update the agents by selecting the **vertical ellipsis › Update** against the agent.

d) Click **Next**.

7. From the **Provide Credentials** page, depending on whether you have a common credential for all VMs managed by the vCenter cloud account or different credentials for all VMs managed by the vCenter cloud account, enter the following details:
   a) If the selected VMs have a common user name and password, enter the common user name and password.
   b) For different user names and passwords for each VM, download the CSV template and add the required details such as the user name, password for each VM managed by the vCenter cloud account. Use the **Browse** button to select the template.
   c) The **Create run time user on Linux virtual machines, with required permissions as part of agent installation** check box is selected by default. For more information, see User Account Prerequisites.
   d) Click **Next**.
8. From the **Summary** page, you can view the list of VMs managed by the vCenter cloud account on which the agent is to be deployed.
9. Click **Install Agent**. Refresh the UI to view the agents that are installed.
   On User Account Control (UAC) deactivated Windows VMs managed by the vCenter cloud account, the Telegraf agent installation completes with the **Agent running** agent status and **Install Success** Last operation status, if all prerequisites are performed.
   **UAC Activated on Windows VMs Managed by the vCenter Server Cloud Account**

   In the case of UAC activated VMs, for a non-administrator user who belongs to the administrator group, Telegraf agent installation completes with the **Not Started** agent status and **Download Success** last operation status, if all the prerequisites are performed. The bits are downloaded to the VM. You must manually install the bits. From `C:\VMware\UCP\downloads`, run a bootstrap launcher, and then the agent status becomes **Agent Running** and the Last operation status becomes **Install Success**.

   - Open PowerShell with administrator privileges.
   - Go to `$SYSTEMDRIVE\VMware\UCP\downloads` folder: `cd $SYSTEMDRIVE\VMware\UCP\downloads`
   - Run the **`cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1`** command.
   - View the results from `uaf_bootstrap.log`.
   - Verify the status of agent installation from the **Agent Status** and **Last Operation Status** columns in the **Manage Telegraf Agents** page.
10. The agent discovers the application services that are installed on the target machines. You can view the application services by clicking the drop down arrow against the VMs managed by the vCenter cloud account on which the agent is installed in the **Manage Telegraf Agents** page. You can view the status of agent installation from the **Agent Status** column in the **Manage Telegraf Agents** page.

11. After installation, an OS object will be created under VMs managed by the vCenter cloud account topology. To view the object, select a VM from **Operations › Applications › Manage Telegraf Agents** page. Click the vertical ellipsis and select **Go to Details**, and click on the **Metrics** tab. Expand the **Object relationship** widget, and select the OS object.
12. If agent installation fails, a detailed error message is displayed when you hover over the red exclamation icon before the Install Failed text in the **Last operation status** column. You can take appropriate action based on the error message to resolve the agent installation issues.

You can manage the services on each agent.

For information about uninstalling an agent, see Uninstall an Agent.

### Install/Uninstall an Agent Using a Script on a Linux Platform

You can install or uninstall an agent on a target machine using a script.

- Complete all the prerequisites. For more information, see Prerequisites.
- The unzip package must be available on the target machine.
- The user must have access permissions to the download folder.
- The cloud account must be configured for the vCenter to which the VM belongs.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account. For more information about object types, see OS and Application Monitoring.

- The VM managed by the vCenter cloud account must be available in VMware Aria Operations.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account.

- The guest IP must be properly configured and is unique across vCenter Servers. If more than one VM is managed by the vCenter cloud account with the same IP is monitored, the script cannot resolve and subscribe to application monitoring.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account. For more information about object types, see OS and Application Monitoring.

- Port 443 in VMware Aria Operations must be accessible to the target machine.
- The VMware Aria Operations user must have the following permissions:
  - Administration > Control Panel > REST-APIs > All other Read, Write APIs
  - Administration >Control Panel > REST-APIs > Read access to APIs
  - Operations > Applications > Manage Telegraf Agent
  - Inventory > Actions > Bootstrap virtual machines
  - Inventory > Actions > Download bootstrap

  **NOTE**
  To check permissions, navigate to **Administration › Control Panel › Access Control › Roles**, select the role to which the user is assigned. You can view the permissions on the right side at the end of the page.

- Only IPv4 is supported at present for cloud proxy.
- VMTools version >=10.2.

  **NOTE**
  Applicable only to VMs managed by the vCenter cloud account.

1. Log in to the target machine on which you want to install/uninstall the agent and download the sample script from cloud proxy from the following location: `https://<CloudProxy>/downloads/salt/telegraf-utils.sh`.

   Run one of the following commands:

   ```
   wget --no-check-certificate https://<CloudProxy>/downloads/salt/telegraf-utils.sh

   curl -k "https://<CloudProxy>/downloads/salt/telegraf-utils.sh" --output telegraf-utils.sh
   ```

   **NOTE**
   Use the relevant cloud proxy **IP address** for <CloudProxy> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2. Make the script executable by running the following command:

   ```
   chmod +x telegraf-utils.sh
   ```

3. To execute the script and install/uninstall the agent, run the following command:

   ```
   telegraf-utils.sh  product-managed -c cloud_proxy_ip_or_collector_group_name -t token (-v vmwareariaoperations_ip_or_fqdn | [-g gateway_url -a csp_auth_url]) [-d download_tmp_dir -s sleep_seconds -i list_of_IPs -o operation]



   Description of arguments:

         -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or Collector Group Name

         -t : [REQUIRED] token - This can be user_saas_refresh_token or on-prem vmwareariaoperations_auth_token.

                saas: CSP Refresh Token of the user/account. For getting new token, follow - User/

                      Organization Settings >> My Account >> API Tokens >> Generate a New API Token.

                      Example: gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpqon-prem: Auth Token of the user/account. For getting new token

                         ( https://<VMwareAriaOperations_IP>/suite-api/ or curl -ks -X POST https://<VMwareAriaOperations_IP>/suite-api/api/auth/token/acquire -H \"Content-Type: application/json\" -H \"Accept: application/json\" -d \"{\"username\": \"<VMwareAriaOperations_USER>\",\"password\": \"<VMwareAriaOperations_USER_PASSWORD>\"}\")

         -d : [OPTIONAL] download_tmp_dir - Temporary directory for agent installation. Default: current directory

         -v : [CONDITIONAL] [ON-PREM-SPECIFIC] vmwareariaoperations_ip_or_fqdn - IP/FQDN of VMware Aria Operations master node and required for on-prem.-g : [OPTIONAL] [SAAS-SPECIFIC] gateway_url - argument to override default VMware Aria Operations SaaS gateway URL

         -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP authentication URL

         -o : [OPTIONAL] operation - The type of operation to be performed {install,uninstall} Default: install
   ```

```
      -s : [OPTIONAL] sleep_seconds - Time (in seconds) to delay the script
execution. This optional parameter will be helpful when this script is used in vRA to
have agent installed on deploying VM. Recommended time 600 seconds.

      -i : [OPTIONAL] list_of_IPs - This is an optional parameter. Comma separated
IPs that are set as properties of a VM in VMware Aria Operations (VM -> Properties ->
Network -> <Integer> -> IP Address). If single adapter has multiple IP Addresses,
then delimit them using '_'

                    ex:- If your VM has 3 adapters with IP Address as follows:
Adapter1: 10.0.0.1 Adapter2: 10.0.0.2 & 10.0.0.3 Adapter3: 10.0.0.4

                    Then this parameter should be given as
"10.0.0.1,10.0.0.2_10.0.0.3,10.0.0.4"


Example: /bin/bash telegraf-utils.sh product-managed -t 8dab02cc-277c-4392-b910-
bd2e98c7e741::8bcde100-6318-44d7-a8dc-11f4ff84b3b -v 10.10.10.100 -c 10.10.10.101
```

**NOTE**
The `-c` argument is mandatory to run the helper script.

To verify the bootstrap status, view the `uaf-bootstrap-results` file. If the installation fails, look for error messages in `uaf_bootstrap.log`.

If the script is successful, the agent status will be updated in the **Manage Telegraf Agents** tab after one collection cycle that takes 5–10 minutes.

**NOTE**
When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

### Install/Uninstall an Agent Using a Script on a Windows Platform

You can install an agent on a target machine using a script.

- Complete all the prerequisites. For more information, see Prerequisites.
- The unzip package must be available on the target machine.
- The user must have access permissions to the download folder.
- Windows PowerShell must be >= 5.0.
- The guest IP must be properly configured and is unique across vCenter Servers. If more than one VM is managed by the vCenter cloud account with the same IP is monitored, the script cannot resolve and subscribe to application monitoring.
- **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account. For more information about object types, see OS and Application Monitoring.
- The VM managed by the vCenter cloud account must be available in VMware Aria Operations.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account.
- The cloud account must be configured for the vCenter to which the VM belongs.

**NOTE**
This prerequisite is applicable only to the VMs managed by the vCenter cloud account.

- Port 443 in VMware Aria Operations must be accessible to the target machine.
- The VMware Aria Operations user must have the following permissions:
  - Administration > Control Panel > REST-APIs > All other Read, Write APIs
  - Administration > Control Panel > REST-APIs > Read access to APIs
  - Operations > Applications > Manage Telegraf Agent
  - Inventory > Actions > Bootstrap virtual machines
  - Inventory > Actions > Download bootstrap

  **NOTE**
  To check permissions navigate to **Administration › Access Control › Roles**, select the role to which the user is assigned. You can view the permissions on the right side at the end of the page.

- Only IPv4 is supported at present for cloud proxy.
- VMTools version must be >=10.2.

  **NOTE**
  Applicable only to VMs managed by the vCenter cloud account.

1. Log in to the target machine on which you want to install/uninstall the agent and download the sample script from cloud proxy from the following location: `https://<CloudProxy>/downloads/salt/ telegraf-utils.ps1`

   If the script download fails with the following message: `The request was aborted: Could not create SSL/TLS secure channel`, follow the steps mentioned in Script Download Fails on a Windows Platform.

   Run one of the following commands:
   `PowerShell command:`

   ```
   Invoke-WebRequest "https://<CloudProxy>/downloads/salt/telegraf-utils.ps1" -OutFile
   telegraf-utils.ps1
   ```

   ```
   Or if you have the wget tool:

   wget --no-check-certificate https://<CloudProxy>/downloads/salt/telegraf-utils.ps1
   ```

   **NOTE**
   Use the relevant cloud proxy **IP address** for <CloudProxy> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2. To execute the script and install/uninstall the agent, run the following command:

   ```
   telegraf-utils.ps1  product-managed -c cloud_proxy_ip_or_collector_group_name -t
   token (-v vmwarearіaoperations_ip_or_fqdn | [-g gateway_url -a csp_auth_url]) [-d
   download_tmp_dir -s sleep_seconds -i list_of_IPs -o operation]
   ```

   ```
   Description of arguments:

         -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or
   Collector Group Name

         -t : [REQUIRED] token - This can be user_saas_refresh_token or on-prem
   vmwarearіaoperations_auth_token.
   ```

```
            saas: CSP Refresh Token of the user/account. For getting new token,
follow - User/

                  Organization Settings >> My Account >> API Tokens >> Generate a New
API Token.

                  Example:
gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpqon-prem: Auth Token
of the user/account. For getting new token

                     ( https://<VMwareAriaOperations_IP>/suite-api/ or curl -ks -X
POST https://<VMwareAriaOperations_IP>/suite-api/api/auth/token/acquire -H \"Content-
Type: application/json\" -H \"Accept: application/json\" -d \"{\"username\":
\"<VMwareAriaOperations_USER>\",\"password\":
\"<VMwareAriaOperations_USER_PASSWORD>\"}\")

      -d : [OPTIONAL] download_tmp_dir - Temporary directory for agent installation.
Default: current directory

      -v : [CONDITIONAL] [ON-PREM-SPECIFIC] vmwarearіaoperations_ip_or_fqdn - IP/FQDN
of VMware Aria Operations master node and required for on-prem.-g : [OPTIONAL] [SAAS-
SPECIFIC] gateway_url - argument to override default VMware Aria Operations SaaS
gateway URL

      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URL

      -o : [OPTIONAL] operation - The type of operation to be performed
{install,uninstall} Default: install

      -s : [OPTIONAL] sleep_seconds - Time (in seconds) to delay the script
execution. This optional parameter will be helpful when this script is used in vRA to
have agent installed on deploying VM. Recommended time 600 seconds.

      -i : [OPTIONAL] list_of_IPs - This is an optional parameter. Comma separated
IPs that are set as properties of a VM in VMware Aria Operations (VM -> Properties ->
Network -> <Integer> -> IP Address). If single adapter has multiple IP Addresses,
then delimit them using '_'

                     ex:- If your VM has 3 adapters with IP Address as follows:
Adapter1: 10.0.0.1 Adapter2: 10.0.0.2 & 10.0.0.3 Adapter3: 10.0.0.4

                     Then this parameter should be given as
"10.0.0.1,10.0.0.2_10.0.0.3,10.0.0.4"



Example: .\telegraf-utils.ps1 product-managed -t 8dab02cc-277c-4383-b910-
bd2e89c7e741::8bcde100-6318-44d7-a8dc-11f4ff84b3b -v 10.10.10.100 -c 10.10.10.101
```

**NOTE**
The `-c` argument is mandatory to run the helper script.

**NOTE**
Do not use a space in the configuration path. Paths with spaces can be passed as a short name notation, such as `c:\PROGRA~1` for `c:\Program Files`.

To verify the bootstrap status, view the `uaf-bootstrap-results` file. If the installation fails, look for error messages in `uaf_bootstrap.log`.

If the script is successful, the agent status will be updated in the **Manage Telegraf Agents** tab after one collection cycle that takes 5–10 minutes.

**NOTE**
When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

### Install/Uninstall an Agent Using a Script on Linux Platforms

You can install or uninstall an agent on a target machine using a script.

- Complete all the prerequisites. For more information, see Prerequisites.
- The unzip package must be available on the target machine.
- The user must have access permissions to the download folder.
- The cloud account must be configured for the vCenter to which the VM belongs. The application monitoring adapter that is mapped to the vCenter is created if it does not exist.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account. For more information about object types, see OS and Application Monitoring.

- The VM managed by the vCenter cloud account must be available in VMware Cloud Foundation Operations.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account.

- The guest IP must be properly configured and should be unique across vCenter Servers. If more than one VM is managed by the vCenter cloud account with the same IP is monitored, the script cannot resolve and subscribe to application monitoring.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account. For more information about object types, see OS and Application Monitoring.

- Ensure that the Internet is enabled.
- The VMware Aria Operations user must have the following permissions:
  - Administration > Control Panel > REST-APIs > All other Read, Write APIs
  - Administration > Control Panel > REST-APIs > Read access to APIs
  - Operations > Applications > Manage Telegraf Agent
  - Inventory > Actions > Bootstrap virtual machines
  - Inventory > Actions > Download bootstrap

  **NOTE**
  To check permissions, navigate to **Administration › Control Panel › Access Control › Roles**, select the role to which the user is assigned. You can view the permissions on the right side at the end of the page.

- Only IPv4 is supported at present for cloud proxy.
- VMTools version must be >=10.2.

  **NOTE**
  Applicable only to VMs managed by the vCenter cloud account.

1. Log in to the target machine on which you want to install/uninstall the agent and download the sample script from cloud proxy from the following location: `https://<CloudProxy>/downloads/salt/ telegraf-utils.sh`. Run one of the following commands:

```
wget --no-check-certificate “https://<CloudProxy>/downloads/salt/telegraf-utils.sh”

curl -k “https://<CloudProxy>/downloads/salt/telegraf-utils.sh” --output telegraf-utils.sh
```

**NOTE**
Use the relevant cloud proxy **IP address** for <CloudProxy> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2. Make the script executable by running the following command:

```
chmod +x telegraf-utils.sh
```

3. Create the user's refresh token associated with the current organization from the following location in the Cloud Service portal: `User/Organization Settings > My Account > API Tokens > Generate a New API Token`
4. To execute the script and install/uninstall the agent, run the following command:

```
telegraf-utils.sh  product-managed -c cloud_proxy_ip_or_collector_group_name -t token (-v vmwareariaoperations_ip_or_fqdn | [-g gateway_url -a csp_auth_url]) [-d download_tmp_dir -s sleep_seconds -i list_of_IPs -o operation]



Description of arguments:

      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or Collector Group Name

      -t : [REQUIRED] token - This can be user_saas_refresh_token or on-prem vmwareariaoperations_auth_token.

            saas: CSP Refresh Token of the user/account. For getting new token, follow - User/

                  Organization Settings >> My Account >> API Tokens >> Generate a New API Token.

                  Example:
gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpqon-prem: Auth Token of the user/account. For getting new token

                     ( https://<VMwareAriaOperations_IP>/suite-api/ or curl -ks -X POST https://<VMwareAriaOperations_IP>/suite-api/api/auth/token/acquire -H \"Content-Type: application/json\" -H \"Accept: application/json\" -d \"{\"username\": \"<VMwareAriaOperations_USER>\",\"password\": \"<VMwareAriaOperations_USER_PASSWORD>\"}\")

      -d : [OPTIONAL] download_tmp_dir - Temporary directory for agent installation. Default: current directory

      -v : [CONDITIONAL] [ON-PREM-SPECIFIC] vmwareariaoperations_ip_or_fqdn - IP/FQDN of VMware Aria Operations master node and required for on-prem. -g : [OPTIONAL] [SAAS-SPECIFIC] gateway_url - argument to override default VMware Aria Operations SaaS gateway URL

      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP authentication URL

      -o : [OPTIONAL] operation - The type of operation to be performed {install,uninstall} Default: install
```

```
      -s : [OPTIONAL] sleep_seconds - Time (in seconds) to delay the script
execution. This optional parameter will be helpful when this script is used in vRA to
have agent installed on deploying VM. Recommended time 600 seconds.

      -i : [OPTIONAL] list_of_IPs - This is an optional parameter. Comma separated
IPs that are set as properties of a VM in VMware Aria Operations (VM -> Properties ->
Network -> <Integer> -> IP Address). If single adapter has multiple IP Addresses,
then delimit them using '_'

                    ex:- If your VM has 3 adapters with IP Address as follows:
Adapter1: 10.0.0.1 Adapter2: 10.0.0.2 & 10.0.0.3 Adapter3: 10.0.0.4

                    Then this parameter should be given as
"10.0.0.1,10.0.0.2_10.0.0.3,10.0.0.4"


Example: /bin/bash telegraf-utils.sh product-managed -t 8dab02cc-277c-4383-b910-
bd2e89c7e741::8bcde100-6318-44d7-a8dc-11f4ff84b3b -v 10.10.10.100 -c 10.10.10.101
```

**NOTE**
The `-c` argument is mandatory to run the helper script.

**NOTE**
The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

To verify the bootstrap status, view the `uaf-bootstrap-results` file. If the installation fails, look for error messages in `uaf_bootstrap.log`.

If the script is successful, the agent status will be updated in the **Manage Telegraf Agents** tab after one collection cycle that takes 5–10 minutes.

**NOTE**
When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

### Install/Uninstall an Agent Using a Script on a Windows Platform

You can install or uninstall an agent on a target machine using a script.

- Complete all the prerequisites. For more information, see Prerequisites.
- The unzip package must be available on the target machine.
- The user must have access permissions to the download folder.
- The cloud account must be configured for the vCenter to which the VM belongs. The application monitoring adapter that is mapped to the vCenter is created if it does not exist.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account. For more information about object types, see OS and Application Monitoring.

- The VM managed by the vCenter cloud account must be available in VMware Cloud Foundation Operations.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account.

- The guest IP must be properly configured and should be unique across vCenter Servers. If more than one VM is managed by the vCenter cloud account with the same IP is monitored, the script cannot resolve and subscribe to application monitoring.

  **NOTE**
  This prerequisite is applicable only to the VMs managed by the vCenter cloud account. For more information about object types, see OS and Application Monitoring.

- Ensure that the Internet is enabled.
- The VMware Cloud Foundation Operations user must have the following permissions:
  - Administration > REST-APIs > All other Read, Write APIs
  - Administration > REST-APIs > Read access to APIs
  - Environment > Applications > Manage Telegraf Agent
  - Environment > Actions > Bootstrap virtual machines
  - Environment > Actions > Download bootstrap

  **NOTE**
  To check permissions navigate to **Administration › Access Control › Roles**, select the role to which the user is assigned. You can view the permissions on the right side at the end of the page.

- Only IPv4 is supported at present for cloud proxy.
- VMTools version must be >=10.2.

  **NOTE**
  Applicable only to VMs managed by the vCenter cloud account.

1. Log in to the target machine on which you want to install/uninstall the agent, download the sample script from cloud proxy from the following location: `https://<CloudProxy>/downloads/salt/telegraf-utils.ps1 .`

   If the script download fails with the following message: `The request was aborted: Could not create SSL/TLS secure channel`, follow the steps mentioned in Script Download Fails on a Windows Platform.

   Run one of the following commands:
   `PowerShell command:`

   ```
   Invoke-WebRequest "https://<CloudProxy>/downloads/salt/telegraf-utils.ps1" -OutFile telegraf-utils.ps1
   ```

   `Or if you have the wget tool:`

   ```
   wget --no-check-certificate https://<CloudProxy>/downloads/salt/telegraf-utils.ps1
   ```

   **NOTE**
   Use the relevant cloud proxy **IP address** for <CloudProxy> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2. Create the user's refresh token associated with the current organization from the following location in the Cloud Service portal: `User/Organization Settings > My Account > API Tokens >> Generate a New API Token`
3. To execute the script and install/uninstall the agent, run the following command:

   ```
   telegraf-utils.ps1  product-managed -c cloud_proxy_ip_or_collector_group_name -t token (-v vmwarearіaoperations_ip_or_fqdn | [-g gateway_url -a csp_auth_url]) [-d download_tmp_dir -s sleep_seconds -i list_of_IPs -o operation]
   ```

```
Description of arguments:

      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or
Collector Group Name

      -t : [REQUIRED] token - This can be user_saas_refresh_token or on-prem
vmwarearia operations_auth_token.

            saas: CSP Refresh Token of the user/account. For getting new token,
follow - User/

                  Organization Settings >> My Account >> API Tokens >> Generate a New
API Token.

                  Example:
gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpqon-prem: Auth Token
of the user/account. For getting new token

                     ( https://<VMwareAriaOperations_IP>/suite-api/ or curl -ks -X
POST https://<VMwareAriaOperations_IP>/suite-api/api/auth/token/acquire -H \"Content-
Type: application/json\" -H \"Accept: application/json\" -d \"{\"username\":
\"<VMwareAriaOperations_USER>\",\"password\":
\"<VMwareAriaOperations_USER_PASSWORD>\"}\")

      -d : [OPTIONAL] download_tmp_dir - Temporary directory for agent installation.
Default: current directory

      -v : [CONDITIONAL] [ON-PREM-SPECIFIC] vmwarearia operations_ip_or_fqdn - IP/FQDN
of VMware Aria Operations master node and required for on-prem.-g : [OPTIONAL] [SAAS-
SPECIFIC] gateway_url - argument to override default VMware Aria Operations SaaS
gateway URL

      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URL

      -o : [OPTIONAL] operation - The type of operation to be performed
{install,uninstall} Default: install

      -s : [OPTIONAL] sleep_seconds - Time (in seconds) to delay the script
execution. This optional parameter will be helpful when this script is used in vRA to
have agent installed on deploying VM. Recommended time 600 seconds.

      -i : [OPTIONAL] list_of_IPs - This is an optional parameter. Comma separated
IPs that are set as properties of a VM in VMware Aria Operations (VM -> Properties ->
Network -> <Integer> -> IP Address). If single adapter has multiple IP Addresses,
then delimit them using '_'

                     ex:- If your VM has 3 adapters with IP Address as follows:
Adapter1: 10.0.0.1 Adapter2: 10.0.0.2 & 10.0.0.3 Adapter3: 10.0.0.4

                     Then this parameter should be given as
"10.0.0.1,10.0.0.2_10.0.0.3,10.0.0.4"



Example: .\telegraf-utils.ps1 product-managed -t 8dab02cc-265c-4392-b910-
bd2e89c7e741::8bcde100-6318-44d7-a8dc-11f4ff84b3b -v 10.10.10.100 -c 10.10.10.101
```

**NOTE**
The `-c` argument is mandatory to run the helper script.

**NOTE**
The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

**NOTE**
Do not use a space in the configuration path. Paths with spaces can be passed as a short name notation, such as `c:\PROGRA~1` for `c:\Program Files`.

To verify the bootstrap status, view the `uaf-bootstrap-results` file. If the installation fails, look for error messages in `uaf_bootstrap.log`.

If the script is successful, the agent status will be updated in the **Manage Telegraf Agents** tab after one collection cycle that takes 5–10 minutes.

**NOTE**
When you use an automation script, concurrent agent installation with a batch size of 20 is supported.

### Uninstall an Agent

You must select the VMs managed by the vCenter cloud account on which you want to uninstall the agent.

**NOTE**
You cannot uninstall an agent from the UI for physical servers and VMs unmanaged by the vCenter cloud account (`Endpoint` object type in VMware Aria OperationsVMware Cloud Foundation Operations).

Ensure that you have completed all the prerequisites. For more information, see Prerequisites.

1. From the left menu, select **Operations › Applications**. From the **Applications** panel, select **Manage Telegraf Agents**. You see the **Manage Telegraf Agents** page.
2. From the **Manage Telegraf Agents** page, select the VMs that are managed by the vCenter cloud account on which you want to uninstall the agent.
3. From the **Manage Telegraf Agents** page, click the horizontal ellipsis, and then click **Uninstall**. You see the **Manage Agent** dialog box.
4. From the **How do you want to provide VM Credentials** page, complete the following steps:
   a) If you have a common user name and password for all the VMs, select the **Common username and password** option.
   b) If you have different user names and passwords for all the VMs, select the **Enter virtual machine credentials** option.
   c) Under the section called **Telegraf Configurations**, select the **Telegraf agent configurations - Save a backup of configurations** option if you want to backup active agent configurations available on the VM. You can restore these agent configurations during install of an agent. For more information, see Install an Agent from the UI.
   d) Click **Next**.
5. From the **Provide Credentials** page, depending on whether you have a common credential for all the VMs or different credentials for all VMs, enter the following details:
   a) If your VM has a single user name and password, enter the common user name and password.
   b) For multiple user names and passwords for each VM, download the CSV template and add the details. Use the **Browse** button to select the template.
   c) Click **Next**.

6. From the **Summary** page, you can view the list of VMs on which the agent is uninstalled.
7. Click **Uninstall Agent**. Refresh the UI to view the progress of agent uninstallation.
   The **Agent Status** column and the missing drop down arrow against the VM in the workspace, indicate that uninstallation is complete and that there are no application services discovered on each agent.
   **UAC Activated on Windows Machine**

   The bits are downloaded to the VM. You have to manually uninstall the bits.
   - Open PowerShell with administrator privileges.
   - Go to `$SYSTEMDRIVE\VMware\UCP\downloads` folder: `cd $SYSTEMDRIVE\VMware\UCP\downloads`
   - Run the **`cmd /c uaf-bootstrap-launcher.bat > uaf_bootstrap.log 2>&1`** command.
   - View the results from `uaf_bootstrap.log`.
   - Verify the status of agent installation from the **Agent Status** and **Last Operation Status** columns in the **Manage Telegraf Agents** page.

### Activate an Application Service

To monitor application services running on the target machines, you must configure Telegraf agents in VMware Aria OperationsVMware Cloud Foundation Operations.

> **NOTE**
> The maximum allowed application services count per target machine is 200.

After you have installed the agent, you can activate plugins to monitor application services. Additionally, you can monitor a discovered service on VMs managed and unmanaged by the vCenter cloud account and physical servers.

You can configure discovered services and carry out custom monitoring actions. For custom monitoring actions, see Custom Script, Activate Remote Checks, Monitor Windows Services, and Monitor Linux Processes.

**Prerequisite**

- If plugin activation requires the location of a file (for example, client certificates for SSL Trust) on the target machine, the location and the files should have appropriate read permissions for the *`arcuser`* to access those files.
  Configuration of application services and custom monitoring using Telegraf is referred to as plugins.

  > **NOTE**
  > If the plugin displays a permission denied status, provide the *`arcuser`* with permissions to the file locations that you have specified during plugin activation.

- Linux process activation for Pid files works only if the Pid file and its parent directories have read permission for **Others**.

### Activate an Application Service on a Single Target Machine

To monitor an application service on a single target machine, complete the following steps:

1. From the left menu, click **Operations › Applications**. From the **Applications** panel, click **Manage Telegraf Agents**.
2. Select the target machine on which the agent is running. You can use the filter functionality too. For example, filter by **Agent Status › Agent Running**.
3. Expand the drop-down arrow against the target machine on which the agent is installed. You see the **Services Discovered** section.
4. From the **Services Discovered** section, select a discovered service or custom monitored service, click the vertical ellipsis and then click **Add**.
5. Activate the application service from dialog box that is displayed on the right side.
6. Enter the details for each instance that you add and click **Save**. For configuration details of each application, see Configuring Supported Application Services.
   Fields with a star are mandatory.

For more information about the status details that appear against the application services, see the table called Status Details in Additional Operations from the Manage Telegraf Agents Page .

To edit or delete instances of application services, click the **Edit** or **Delete** options from the vertical ellipsis against application service you added. After the services have been added and saved, click the drop down arrow against the application service to view the list of services and their status.

**Activate an Application Service on Multiple Target Machines**

To monitor an application service on multiple target machines, complete the following steps:

1. From the left menu, click **Operations › Applications**. From the **Applications** panel, click **Manage Telegraf Agents**.
2. Select the target machines on which the agent is running where the agent type is `product-managed`. You can use the filter functionality too. For example, filter by **Agent Status › Agent Running** and **Agent Type › Product-managed agent**.

   **NOTE**
   You can only activate an application service for all object types (Endpoint and Virtual Machine) with a product-managed Telegraf agent installed and running.
3. Click the **Horizontal Ellipsis** and then click **Activate Service**. You see the **Activate Service** dialog box.
4. From the **Activate Service** dialog box, select a discovered service or custom monitoring service by clicking the drop-down arrow in the **Select a Service** field and click **Confirm**.
5. Activate the application service from dialog box that is displayed on the right side.
6. Enter the details for each instance that you add and click **Save**. For configuration details of each application, see Configuring Supported Application Services.
   Fields with a star are mandatory.

   For more information about the status details that appear against the application services, see the table called Status Details in Additional Operations from the Manage Telegraf Agents Page .

   To edit or delete instances of application services, click the **Edit** or **Delete** options from the vertical ellipsis against application service you added. After the services have been added and saved, click the drop down arrow against the application service to view the list of services and their status.

   For custom monitoring actions, see Custom Script, Activate Remote Checks, Monitor Windows Services, and Monitor Linux Processes.

The following special characters are permitted in the DB user field: `'[]{} (),.<> ?:!|/~@#$%^&*- _ +=`

You can provide DB name lists in the following format `['DBNAME_1', 'DBNAME_2', 'DBNAME_3']` where DBNAME_1, DBNAME_2, DBNAME_3 must not contain quotes such as `'` and `"`.

**Application Availability**

When an application service is activated, the **Application Availability** metric is collected and displays if the application service is running on the target machine or if it is down. `1` indicates that the application service is running on the target machine and `0` indicates that the application service is down. This metric is available for all supported application services except JAVA application service.
For information about deactivating a service, see Deactivate an Application Service.

## Configuring Supported Application Services

Twenty-three application services are supported in VMware Aria OperationsVMware Cloud Foundation Operations. The supported application services are listed here. Some of the application services have mandatory properties which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected.

Active Directory
Active Directory is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

Active MQ
ActiveMQ is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Server URL | Yes | http://localhost:8161 |
| User name | Yes | User name for Active MQ. Example: admin |
| Password | Yes | Password |
| Installed Path | Yes | The path on the target machine where Active MQ is installed.<br>Example:<br>For Linux target machines: `/opt/apache-activemq`<br>For Windows target machines: `C:\apache-activemq-5.15.2` |

Apache HTTPD
Apache HTTPD is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Status Page URL | Yes | http://localhost/server-status?auto |
| User name | No | User name for Apache HTTPD service. `Example:root` |
| Password | No | Password |
| Response String Match | No | Optional substring or regex match in the body of the response (case sensitive). For example, `ok`. You must escape regex special characters `?,` |

*Table continued on next page*

*Continued from previous page*

| Name | Mandatory? | Comment |
|---|---|---|
| | | `(, ), .` with a backslash. For example, for matching against `?doodles`, escape the `?` with `\?`, such as, `\?doodles`. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

Cassandra Database

Cassandra database is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the application instance. |
| Installed Path | Yes | Valid file path. |
| URL | Yes | http://localhost:8778 |

Hyper-V

Hyper-V is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the application service. |

Java

Java is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |
| Installed Path | Yes | The path on the target machine where Java is installed. Example: For Linux target machines: /opt/vmware/ucp ; For Windows target machines: C:\VMware\UCP |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |

*Table continued on next page*

*Continued from previous page*

| Name | Mandatory? | Comment |
|---|---|---|
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

JBoss

JBoss is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |
| Installed Path | Yes | The path on the target machine where JBoss is installed. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

MongoDB

MongoDB is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where MongoDB is running. Example: 27017 |
| Hostname | No | Optional hostname for the MongoDB Service. |
| Username | No | User name for MongoDB. Example: Root |
| Password | No | Password |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

MS Exchange

MS Exchange is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

MS IIS

MS IIS is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

MS SQL

MS SQL is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Instance | Yes | Instance name of the MS SQL server. |
| Port | No | The port where MS SQL is running. Example: 1433 |
| Hostname | No | Optional hostname for the MS SQL Service. |
| Username | Yes | User name for MS SQL. Example: Root |
| Password | Yes | Password |

MySQL

MySQL is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where MySQL is running. Example: 3306 |
| User name | Yes | User name for MySQL service. Example: Root |
| password | Yes | Password |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Hostname | No | Optional hostname for the MySQL Service |
| Databases | No | Comma-separated list of databases to monitor. Each of the database names |

*Table continued on next page*

*Continued from previous page*

| Name | Mandatory? | Comment |
|---|---|---|
| | | to be monitored must be enclosed in single quotes and the databases themselves should be comma separated. For example, 'database1','database2','database3'. |
| TLS Connection | No | Allowed values are true, false, and skip-verify. |

Nginx

Nginx is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

**Table 55:**

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the application instance. |
| Status Page URL | Yes | http://127.0.0.1:8081/nginx_status |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain and host verification. Expected: True/False. |

NTPD

NTPD is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

Oracle Database

Oracle database is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the application instance. |
| OracleDB Username | Yes | User name for the Oracle database instance. |
| OracleDB Password | Yes | Password for the Oracle database instance. |
| OracleDB SID | Yes | SID of the Oracle database instance. |

Pivotal Server

Pivotal Server is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |
| Installed Path | Yes | The path on the target machine where Pivotal server is installed. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

Postgres

Postgres is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where PostgreSQL is running. Example: 5432 |
| User name | Yes | User name for PostgreSQL service. Example: Root |
| Password | Yes | Password |
| SSL Connection | No | Allowed values are deactivate, verify-ca, verify-full. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: true/false. |
| Hostname | No | Optional hostname for the PostgreSQL Service. |
| Default Database | No | The database for initiating connection with the server |
| Databases | No | Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma-separated, for example , 'database1','database2','database3'. |

*Table continued on next page*

*Continued from previous page*

| Name | Mandatory? | Comment |
|---|---|---|
| Ignored Databases | No | Comma-separated list of databases that need not be monitored. Each of the database names to be excluded from monitoring must be enclosed in single quotes and the databases themselves should be comma-separated for example, 'database1','database2','database3'. |

RabbitMQ

RabbitMQ is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Management Plugin URL | Yes | http://localhost:15672 |
| User name | No | User name for RabbitMQ. Example: Guest |
| Password | No | Password |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |
| Nodes | No | Each of the RabbitMQ data collection nodes should be in single quotes and the nodes themselves should be comma-separated. The list of nodes must be enclosed in square brackets. For example ['rabbit@node1','rabbit@node2',.....] |

Riak

Riak is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Server URL | Yes | http://localhost:8098 |

Sharepoint

Sharepoint is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

Tomcat

Tomcat is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:8080 |
| Installed Path | Yes | The path on the target machine where Tomcat is installed. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on thetarget machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

Weblogic

Weblogic is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Base URL | Yes | http://localhost:7001 |
| Installed Path | Yes | The path on the target machine where WebLogic is installed. |
| User name | Yes | User name for WebLogic. Example: admin |
| Password | Yes | Password |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificate | No | Path to the SSL Certificate file on the target machine. |
| SSL Key | No | Path to the SSL Key file on the target machine. |
| Skip SSL Verification | No | Use SSL but skip chain & host verification. Expected: True/False. |

Websphere

Websphere is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| IBM Websphere Server URL | Yes | Example : http://localhost:9081 |
| Websphere Authorization Token | Yes | To generate the token, follow the below steps:<br>• Go to https://www.base64encode.org.<br>• Type in the user and password created in the format: user:password<br>• Click the Encode button.<br>• Copy the resulting Base64 encoded string. Example: d2F2ZWZyb250OndhdmVmcm9udA== |

Remote Checks
HTTP Remote Check
HTTP is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the remote check instance. |
| URL | Yes | http://localhost |
| Method | Yes | GET/POST/PUT |
| Proxy | No | Proxy URL: http://localhost |
| Response Timeout | No | Timeout for the connection in seconds. For example, 10. |
| Follow Redirects | No | True/False if redirects from the server. For example, true/false (all small values). |
| Body | No | HTTP request body. |
| Response String Match | No | Substring or regex match in the response body. |
| SSL CA | No | Path to the SSL CA file on the target machine. |
| SSL Certificates | No | Path to the SSL certificate file on the target machine. |
| SSL Key | No | Path to the SSL key file on the target machine. |
| Skip Host & chain verification | No | Use SSL but skip chain and host verification. Expected: True/False. |
| Headers | No | HTTP Request Headers, which can look like the following. For example: **accept = "application/json" Authorization = "5609fe9d-cddd-4654-898b-26d9b67f137a:** |

*Table continued on next page*

*Continued from previous page*

| Name | Mandatory? | Comment |
|---|---|---|
| | | `:4ca86c16-ac01-4e17-90da-fb55eb36571a"` |

ICMP Remote Check

ICMP is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the remote check instance. |
| FQDN/IP | Yes | Host name to send the packets. Example: *example.org* |
| Count | No | Number of ping packets to send per interval. For example, 1. |
| Ping Interval | No | Time to wait between ping packets in seconds. For example, 10.0.<br>**NOTE**<br>Follow the decimals as mentioned in the example. |
| Timeout | No | Timeout to wait for ping response in seconds. For example, 10.0.<br>**NOTE**<br>Follow the decimals as mentioned in the example. |
| Deadline | No | The total ping deadline in seconds. For example, 30. |
| Interface | No | Interface or source from which to send a ping. |

TCP Remote Check

TCP is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the remote check instance. |
| Address | Yes | <hostname>:port |
| Send | No | The given string is sent to the TCP. It can be any string of your choice. |
| Expect | No | The given string is expected from the TCP. It can be any string of your choice. |
| Timeout | No | Timeout for the connection to the TCP server. For example, 10. |
| Read Timeout | No | Timeout for the response from the TCP server. For example, 10. |

UDP Remote Check

UDP is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display name of the remote check instance. |
| Address | Yes | <hostname>:port |
| Send | Yes | The given string is sent to the UDP. |
| Expect | Yes | The given string is expected from the UDP. |
| Timeout | No | Timeout for the connection to the UDP server. For example, 10. |
| Read Timeout | No | Timeout for the response from the UDP server. For example, 10. |

## Configuring Supported VeloCloud Services

Eight VeloCloud application services are supported in VMware Aria OperationsVMware Cloud Foundation Operations. The supported application services are listed here. Some of the application services have mandatory properties which you must configure. Some of the application services have pre-requirements that you must configure first. After you configure the properties, data is collected.

VeloCloud Orchestrator

VeloCloud Orchestrator and the following services are supported in VMware Aria OperationsVMware Cloud Foundation Operations.

- VeloCloud Orchestrator
- Nginx

  **NOTE**
  To activate the plugin for ngnix service you must use the loopback address in the url `http://127.0.0.1/nginx_status`.

- Clickhouse
- Network Time Protocol
- MySQL
- Redis
- Java Application

  **NOTE**
  Java application gets discovered after bootstrapping a VeloCloud Orchestrator virtual machine, but you must ignore it, as we do not monitor the Java application.

In VeloCloud Orchestrator, we monitor the following services. For each of these services we display a metric which indicates the service status:

- Backend
- Portal
- Upload

VeloCloud Orchestrator details.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the VeloCloud Orchestrator instance. |

Nginx

Nginx is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Status Page URL | Yes | http://127.0.0.1/nginx_status |
| SSL CA | No | Path to the SSL CA file on the end point VM. |
| SSL Certificate | No | Path to the SSL Certificate file on the end point VM. |
| SSL Key | No | Path to the SSL Key file on the end point VM. |
| Skip SSL Verification. | No | Use SSL but skip chain & host verification. Expected: True/False. |

ClickHouse

ClickHouse is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Servers URL | Yes | http://127.0.0.1:8123 |
| User name | No | User name for the ClickHouse service. |
| Password | No | Password |

NTPD

NTPD is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

MySQL

MySQL is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

To activate the MySQL plug-in and fetch the credentials, refer to the article Steps to fetch password for telegraf user of MySQL, while activating plugin (81153) at the VMware Support Knowledge Base.

Use the port number 3306 to run MySQL and the telegraf credentials and activate the plug-in.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Port | Yes | The port where MySQL is running. Example: 3306 |
| User name | Yes | User name for the MySQL service. Example: Root |
| password | Yes | Password |
| SSL CA | No | Path to the SSL CA file on the end point VM. |
| SSL Certificate | No | Path to the SSL Certificate file on the end point VM. |
| SSL Key | No | Path to the SSL Key file on the end point VM. |
| Hostname | No | Optional hostname for the MySQL Service |
| Databases | No | Comma-separated list of databases to monitor. Each of the database names to be monitored must be enclosed in single quotes and the databases themselves should be comma-separated. For example, 'database1','database2','database3'. |
| TLS Connection | No | Accepted values are true, false, and skip-verify. |

Redis

Redis is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |
| Redis URL | Yes | servers = ["tcp://localhost:6379"] |
| SSL CA | No | Secure Socket Layer Certification Authority. |
| SSL Certificate | No | Secure Socket Layer Certificate. |
| SSL Key | No | Secure Socket Layer Key |
| Skip SSL Verification. | No | Skips verification for SSL. |

VeloCloud Gateway

VeloCloud Gateway and the following services are supported in VMware Aria OperationsVMware Cloud Foundation Operations

- Network Time Protocol
- VeloCloud Gateway

In VeloCloud Gateway, we monitor the following processes. For each of these processes, we display a metric which indicates the process status.

- bgpd
- watchquagga

- gwd
- mgd
- natd
- ssh
- vc procmon
- vcsyscmd

VeloCloud Gateway details.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the VeloCloud Gateway instance. |

NTPD
NTPD is supported in VMware Aria OperationsVMware Cloud Foundation Operations.

| Name | Mandatory? | Comment |
|---|---|---|
| Display Name | Yes | Display Name of the application instance. |

**Pre-Requirements for Application Services**

For the Telegraf agent to collect metrics for some of the application services, you must make modifications in the target machines. After you make these modifications, the agent starts collecting metrics. You must SSH to the virtual machine where you have deployed the agent and modify the configuration files.

Apache HTTPD
For the agent to collect metrics, modify the conf file available in `/etc/httpd/conf.modules.d/status.conf` and activate the mod_status for the HTTPD plugin.

```
<IfModule mod_status.c>

<Location /server-status>

    SetHandler server-status

</Location>

ExtendedStatus On

</IfModule>
```

If the conf file is not available, you must create one. Restart the HTTPD service after modifying the conf file with the following command:

```
systemctl restart httpd
```

Java Application

Support for Java application monitoring is currently limited to applications that support jolokia based JMX monitoring. Currently, as part of Java plugin activation, both the jolokia.jar and jolokia.war files are placed in the "Installed Path". If the Java application to be monitored is a servlet container, then the jolokia.war file is automatically deployed to the "Installed Path" that you provided when you configured the Java application in VMware Aria OperationsVMware Cloud Foundation Operations. If auto deployment is not supported or not activated, restart the application after activating the plugin or refer to the application's product documentation for the deployment of manual .war.

For all other java applications, first activate the plugin then include the path to the jolokia.jar file in the `"-javaagent:" JVM flag as a command line parameter`. For example, `"-javaagent:/<path_to_jolokia>/jolokia.jar"`, `<path_to_jolokia>` is the "Installed Path" which was provided during application configuration. Restart the application.

> **NOTE**
> For monitoring a customized third-party java application, ensure that:
> - The jolokia based JMX monitoring is supported.
> - In case auto deploy is turned off, manual .war deployment is supported for servlet container application.
> - `"-javaagent:" JVM flag` is added as part of the command line parameter (JVM_OPTS)

Nginx

Add the following lines to the conf file available in `/etc/nginx/nginx.conf`:

```
http {
    server {
        location /status {
            stub_status on;
   access_log off;
   allow all;
  }
    }
}
```

Restart the Nginx service with the following command:

```
systemctl restart nginx
```

Postgres

In the configuration file available in the `/var/lib/pgsql/data/pg_hba.conf`, change the value of `local all postgres peer` to `local all postgres md5` and restart the service with the following command:

```
sudo service postgresql restart
```

Cassandra

To monitor the Cassandra database application, the Jolokia jar must be included as a JVM input to the Cassandra database application. Complete the following steps:

1. Modify `/etc/default/cassandra`.

```
echo "export JVM_EXTRA_OPTS=\"-javaagent:/usr/share/java/jolokia-jvm-1.6.0-
agent.jar=port=8778,host=localhost\"" | sudo tee -a /etc/default/cassandra
```

**NOTE**
You can download the latest version of the Jolokia agent from https://jolokia.org/download.html. For a .JAR file deployment, you have to restart the application service after including the full file path of the JAR in the JMX argument of the JAVA process which you are monitoring.

2. Alternatively, you can activate the agent by modifying `cassandra-env.sh`. Include the following line at the end of the `cassandra-env.sh`:

   ```
   JVM_OPTS="$JVM_OPTS -javaagent:/usr/share/java/jolokia-jvm-1.6.0-
   agent.jar=port=8778,host=localhost"
   ```

   After you see the JVM inputs, restart the Cassandra service.

Oracle Database

To monitor the Oracle database, complete these steps:

1. Download the instant client library from: `https://www.oracle.com/database/technologies/instant-client/downloads.html`.

   You must download the Oracle instant library and included it in the PATH.
2. Install Python 3.6 or later. Install the cx_Oracle module.
   `python3 -m pip install cx_Oracle --upgrade`
3. Create a User.

   ```
   CREATE USER <UserName> IDENTIFIED BY <yourpassword>;

    GRANT select_catalog_role TO <UserName>;

    GRANT CREATE SESSION TO <UserName>;
   ```
4. **For a Windows VM:**
   a. Append the Oracle instant client library path and the Python path to the PATH environment variable.
   b. Install the ARC agent.
5. **For a Linux VM:**
   a. Install the ARC agent.
   b. Get the file `exec_oracle_python.sh(wget --no-check-certificate https://<CP_IP>/downloads/salt/content-accessories/exec_oracle_python.sh)` from `/opt/vmware/ucp/`.
   c. Fill in `/opt/vmware/ucp/exec_oracle_python.sh` with absolute values for `LD_LIBRARY_PATH`, `ORACLE_HOME`, `PYTHON_BIN`, and `TNS_ADMIN`.
      `LD_LIBRARY_PATH`: Path to the Oracle instant client library. For example, `/opt/vmware/ucp/instantclient_21_4`.

      `ORACLE_HOME`: Directory in the file system where the Oracle software is installed. For example, `/u01/app/oracle/product/19.3.0/dbhome_1/`.

      `PYTHON_BIN`: Path to the Python executable. For example, `/usr/bin/python`.

      `TNS_ADMIN`: Environment variable that points to the directory where the SQL*Net configuration files are located. For example, `/u01/app/oracle/product/19.3.0/dbhome_1/network/admin`.
   d. Activate the Oracle DB plugin.

Active MQ 5.16 and Later Versions

To activate Active MQ 5.16 and later versions, complete the following steps:

- Navigate to `/opt/activemq/apache-activemq-5.16.0/webapps/api/WEB-INF/classes/jolokia-access.xml`
- Remove or comment out the following lines:

```
<cors>
  <strict-checking/>
</cors>
```

- Restart the Active MQ service.

Microsoft SQL Server

The user account must have the following permissions to monitor the Microsoft SQL Server application service with Telegraf.

```
USE master;
GO
CREATE LOGIN [telegraf] WITH PASSWORD = N'mystrongpassword';
GO
GRANT VIEW SERVER STATE TO [telegraf];
GO
GRANT VIEW ANY DEFINITION TO [telegraf];
GO
```

> **NOTE**
> If you want to monitor named instances of the Microsoft SQL Server application service, then the plugin activation input should not have the Port field configured. Use Port only when there is no instance in the server (default 1433).

### Additional Operations from the Manage Telegraf Agents Page

After you have configured cloud proxy and installed a product-managed Telegraf agent, configure a vCenter Cloud account (relevant only to VMs managed by the vCenter cloud account). You can manage the agents on all the target machines from the **Manage Telegraf Agents** page.

From the **Manage Telegraf Agents** page, you can:

- Use the filter to easily find physical servers or VMs managed or unmanaged by the vCenter cloud accounts.
- View data collected on all target machines.
- Start, stop, and update agents on all target machines where the product-managed agent is installed. You can also discover and manage the services on each product-managed Telegraf agent that you install.

> **NOTE**
> The maximum allowed application services count per target machine is 200.

Additionally, all the types of supported monitored object types such as `Virtual Machine` and `Endpoint` are consolidated in the **Manage Telegraf Agents** page so that you have visibility of the entire environment you are monitoring

and can manage agents and monitor the content from one single page, depending on the agent you are using and type of the objects available. For more information about object types, see OS and Application Monitoring.

### Telegraf Agents in VMware Aria OperationsVMware Cloud Foundation Operations

For more information about Telegraf agents, see What Are The Different Types of Telegraf Agents in VMware Aria OperationsVMware Cloud Foundation Operations?.

> **NOTE**
> Data collected from product-managed agents, open-source agents, and user-managed agents can be viewed from the **Manage Telegraf Agents** page. For objects that have open-source agents and user-managed Telegraf agents installed and running, you can only view data that is collected, you cannot install, uninstall, or manage the lifecycle of the agents.

### Where You Manage the Agents

To manage the agents and application services, from the left menu, click **Operations › Applications**. From the **Applications** panel, click the **Manage Telegraf Agents** tab.

> **NOTE**
> In the table below, **Vertical Ellipsis** is applicable to single VM operations and **Agent Actions** is applicable to group actions.

**Table 56: Options**

| Options | Description |
|---|---|
| **Agent Actions/Vertical Ellipsis › Install** | Installs the agents on the selected VM/s managed by the vCenter cloud account. Select the VM/s managed by the vCenter cloud account on which you want to install the agent, click Agent Actions/vertical ellipsis and then click **Install**. For more information, see Install an Agent from the UI.<br>**NOTE**<br>• **Install** from the UI is available only for VMs managed by the vCenter cloud account using the product-managed agent.<br>• For `Endpoint` object types, product-managed agent install is not supported.<br>• **Install** from the UI is not available for open-source agents and user-managed agents. |
| **Agent Actions/Vertical Ellipsis › Uninstall** | Uninstalls the agent. Select the VM/s managed by the vCenter cloud account on which you want to uninstall the agent, click Agent Actions/vertical ellipsis and then click **Uninstall**. For more information, see Uninstall an Agent. |

*Table continued on next page*

*Continued from previous page*

| Options | Description |
|---|---|
| | **NOTE**<br>• Uninstall from the UI is available only for VMs managed by the vCenter cloud account using the product-managed agent.<br>• For `Endpoint` objects project-managed agent uninstall is not supported.<br>• **Uninstall** from the UI is not available for open-source agents and user-managed agents. |
| **Agent Actions Ellipsis/Vertical Ellipsis › Update** | Updates agents that are at a lower version. Select the VMs managed and unmanaged by the vCenter cloud account and physical servers on which you want to update the agent, click Agent Actions/vertical ellipsis and then click **Update**. After the agents are updated, the last operation status changes to **Content Upgrade Success**.<br>**NOTE**<br>You can use the **Update** option only for objects with a product-managed Telegraf agent installed and running.<br>**NOTE**<br>After you upgrade VMware Aria OperationsVMware Cloud Foundation Operations and then update product-managed Telegraf agents that are running at an earlier version, there could be a data collection gap of a maximum of 2-3 collection cycles. |
| **Agent Actions/Vertical Ellipsis › Start** | If you have temporarily stopped sending metrics to VMware Aria OperationsVMware Cloud Foundation Operations, you can use the **Start** option to start data collection for the application service. Select the VMs managed and unmanaged by the vCenter cloud account and physical servers on which you want to start the agent and click Agent Actions/vertical ellipsis and then click **Start**.<br>**NOTE**<br>You can use the **Start** option only for objects with a product-managed Telegraf agent installed and running. |
| **Agent Actions/Vertical Ellipsis › Stop** | During a maintenance period, you can temporarily stop sending application service metrics to VMware Aria OperationsVMware Cloud Foundation Operations. Select the VMs managed and unmanaged by the vCenter cloud account and physical servers on which you want to stop the agent and click Agent Actions/vertical ellipsis and then click **Stop**.<br>**NOTE**<br>You can use the **Stop** option only for objects with a product-managed Telegraf agent installed and running. |
| **Vertical Ellipsis › Go To Details** | Displays the **Summary** tab of the selected VMs managed and unmanaged by the vCenter cloud account and physical servers. |

*Table continued on next page*

*Continued from previous page*

| Options | Description |
|---|---|
| | **NOTE**<br>You can view the summary and details for objects with product-managed Telegraf agents, open-source Telegraf agents, and user-managed Telegraf agents installed and running. |
| **Agent Actions > Activate Services** | Allows you to activate a service.<br>**NOTE**<br>You can only activate a service for objects with a product- managed Telegraf agent installed and running.<br>**NOTE**<br>Applicable for bulk actions. |
| Filter | Filters the VMs managed and unmanaged by the vCenter cloud account and physical servers based on the name of the VM or the host name in the case of `Endpoint` type objects, the operating system it runs on, object type, agent type, agent status, the application service discovered, the last operation status, the power status of the VMs managed by the vCenter cloud account, and the agent version. |

You can also view specific details from the options in the data grid.

**Table 57: Data Grid Options**

| Option | Description |
|---|---|
| Name | Name of the virtual machine or host name in case of physical servers and VMs unmanaged by the vCenter cloud account. |
| Object Type | There are four types of object types:<br>• Virtual Machine: Includes only VMs managed by the vCenter cloud account.<br>• Endpoint: Includes physical servers and VMs unmanaged by the vCenter cloud account.<br>**NOTE**<br>The agent types supported by the `Endpoint` object type are: `product-managed agent`, `open-source agent`, and `user-managed agent`.<br>• EC2 Instance<br>• Azure Virtual Machine<br>**NOTE**<br>EC2 Instance and Azure Virtual Machine object types are listed for open-source Telegraf. |
| Operating System | Operating system running on the target machine. |
| Agent Status | Displays the status of the agent at the target machine.<br>• Agent Running. Indicates that the agent is running. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>Applicable only for product managed Telegraf agents.<br>• Agent Stopped. Indicates that the agent has stopped.<br>• Agent Unhealthy: Indicates that a component of the product-managed agent is not healthy.<br>• Not Installed. Indicates that the agent is not installed. |
| Agent Type | Types of agents:<br>• Product-managed agent: Telegraf agent managed by VMware Aria OperationsVMware Cloud Foundation Operations.<br>• Open-source agent: Open source Telegraf agent.<br>• User-Managed agent: Telegraf agent used to monitor physical servers.<br>**NOTE**<br>The `User-managed agent` option is available only if you have upgraded from a previous version of VMware Aria OperationsVMware Cloud Foundation Operations. |
| Last Operation Status | Status of the last operation.<br>**NOTE**<br>Applicable only for product-managed agents.<br>The possible values are:<br>• No Operation<br>• Install Success<br>• Install Failed<br>• Install In Progress<br>• Start Success<br>• Start Failed<br>• Start In Progress<br>• Stop Success<br>• Stop Failed<br>• Stop In Progress<br>• Update Success<br>• Update Failed<br>• Update In Progress<br>• Uninstall Success<br>• Uninstall Failed<br>• Uninstall In Progress<br>• Download Success |
| State | Power status of the VMs managed by the vCenter cloud account. The possible values are:<br>• Powered On<br>• Powered Off |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>Powered on and Powered Off status is applicable only for the `Virtual Machine` object type.<br>• Unknown |
| Virtual IP Details | Virtual IP configured on an application monitoring high availability activated collector group to which target machines will send data. For collector groups that are not activated for application monitoring high availability, the column will be empty. |
| Collector Group | Displays the name of the collector group to which the cloud proxy mentioned in the **Cloud Proxy** column belongs. |
| Cloud Proxy | You can see either:<br>• The cloud proxy IP through which Telegraf agent management is performed and data is collected, or<br>• For high availability activated collector groups, the IP address of the primary cloud proxy that is collecting data is displayed. |
| Agent Version | Version of the agent. A gray dot is displayed if the target machine requires an update for product managed agents.<br>• For product-managed agents: The version displayed is the same as the VMware Aria OperationsVMware Cloud Foundation Operations version.<br>• For open-source agents: The version displayed is the open source Telegraf agent version that you use. |
| vCenter Name | Name of the vCenter Adapter cloud account to which that vCenter VMs resource belongs.<br>**NOTE**<br>The field is empty for `Endpoint` object types.<br>**NOTE**<br>The field is empty for `Endpoint, EC2 Instance, and Azure VM` object types. |
| Collection State | VMware Aria Operations Application Management cloud account collection state which monitors the VM. |
| Collection Status | VMware Aria Operations Application Management cloud account collection status which monitors the VM. |

**Table 58: Status of Application Services**

| Icon | Description |
|---|---|
| Green tick icon against the application service | Indicates that the application service is activated and the application service instances are receiving data. |
| Red exclamation icon against the application service | Indicates that the application service has been activated but there is a problem with data collection.<br>When there is more than one agent plugin of the same kind, and one of them is activated, but the other is not |

*Table continued on next page*

*Continued from previous page*

| Icon | Description |
|---|---|
| | collecting data, a red exclamation icon is still displayed against the application service. |
| Red icon against the application service instance | Indicates that there is an error in receiving data with an error message displayed after the name of the application service instance. |
| Gray question icon against the application service | Indicates that the application services requires reactivation. The application service must be reactivated. For reactivation, see Activate an Application Service for more information |
| Gray pause icon against the application service instance | Indicates that the application service instance has been stopped. |
| Blue icon with three horizontal dots against the agent plugin | Indicates data is being received. |
| Progress status icon | After you have added the parameters and activated the application service, the progress status is displayed until data collection starts. |

To manage the agent, follow these steps:

1. Install the agent.
   For more information, see Install an Agent from the UI.

   **NOTE**
   - You can install the agent from the UI only for VMs managed by the vCenter cloud account.
   - To install agents on physical servers and VMs unmanaged by the vCenter cloud account, use the script. For information to install/uninstall product-managed Telegraf agents, see Install/Uninstall an Agent Using a Script on a Linux Platformand Install/Uninstall an Agent Using a Script on a Windows PlatformFor more information to install product-managed Telegraf agents, see Install/Uninstall an Agent Using a Script on Linux Platformsand Install/Uninstall an Agent Using a Script on a Windows Platform . For information to install open-source agents, see Monitoring Applications using Open Source Telegraf on a Linux Platform and Monitoring Applications using Open Source Telegraf on a Windows Platform.

2. Manage the application services on each agent.
   For more information, see Activate an Application Service.

   **NOTE**
   You can manage application services only for objects with a product-managed Telegraf agent installed and running.

   For custom monitoring actions, see Custom Script, Activate Remote Checks, Monitor Windows Services, and Monitor Linux Processes.
3. Stop and start the agents on the target machines.
   When you stop an agent, you cannot activate or deactivate a plugin. If the target machine is powered off or if you lose connection with cloud proxy, you cannot configure or activate a plugin.

   **NOTE**
   You can stop and start agents only for objects with a product-managed Telegraf agent installed and running.

4. Uninstall the agent.
   For more information, see Uninstall an Agent.

**NOTE**
You can uninstall the agent from the UI only for managed vCenter VMs.

5. Update agents that are at a lower version.

### What Are The Different Types of Telegraf Agents in VMware Aria OperationsVMware Cloud Foundation Operations ?

The three types of Telegraf agents are product-managed agent, open source agent, and user-managed agent in VMware Aria OperationsVMware Cloud Foundation Operations.

- **Product-managed agent**: Managed by VMware Aria OperationsVMware Cloud Foundation Operations. You can install and uninstall (only for VMs managed by the vCenter Server cloud account), activate application services, and manage the life cycle of the Telegraf agents such as upgrade the agent, start the agent, stop the agent, and so on. For more details, see Additional Operations from the Manage Telegraf Agents Page .
- **Open source agent**: Refers to the open source Telegraf agent that you use to monitor curated and non-curated application services completely by yourself. You can install the open source Telegraf agent using a script and also configure it using the helper script to send data to VMware Aria OperationsVMware Cloud Foundation Operations. For more information, see Monitoring Applications using Open Source Telegraf on a Linux Platform and Monitoring Applications using Open Source Telegraf on a Windows Platform. You cannot manage the life-cycle of open source telegraf agents in VMware Aria OperationsVMware Cloud Foundation Operations, such as activate an application service, and stop/start upgrade. You can only view configured application instances and their data collection. Agent management is carried out on the target machine by the user.
- **User-managed agent**: User-managed agents were used to monitor physical servers in prior releases. Please note the following:
  - Available only if you upgrade from previous versions of VMware Aria OperationsVMware Cloud Foundation Operations. If you want to continue with the existing user-managed agents you can only view the metrics collected and you do not have to make any changes in VMware Aria OperationsVMware Cloud Foundation Operations.
  - Deprecated for fresh install. You will not be able to install new user-managed agents.
  - As a first time user if you want to monitor physical servers, use the product-managed Telegraf agent and use the helper script for installation of the Telegraf agent. For more information, see Install/Uninstall an Agent Using a Script on a Linux Platform and Install/Uninstall an Agent Using a Script on a Windows PlatformInstall/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform

### Custom Script

You can run custom scripts in the target machine and collect custom data which can then be consumed as a metric.

Prerequisites

- All the scripts that you run using the custom script, must output a single integer value. If the output is not a single integer value, an error is displayed in the user interface.
- The custom script uses Telegraf's `exec` plugin to run scripts on a target machine's operating system. In Linux operating systems, a special user called *`arcuser`* with specific privileges, is created for installing the Telegraf agent. As a result, the `exec` plugin runs the scripts using that *`arcuser`* user. Ensure that the *`arcuser`* can run the scripts that use the custom script (the *`arcuser`* must have permissions to run the script). For example, the *`arcuser`* created automatically by cloud proxy, does not have privileges to run scripts which are stored under the `/root` directory.
- In Windows operating systems, a system user is used for installing the Telegraf agent. As a result, the `exec` plugin runs the scripts using that system user. Ensure that the system user has privileges and can run the custom script.
- The script must be placed in the `/opt/vmware` folder.

- The environment variable PATH must be set for PowerShell scripts to be executed on the Windows operating system if the user used the *powershell -File* prefix. For example, *"%SYSTEMROOT%\System32\WindowsPowerShell\v1.0\"* in the environment variable PATH variable.

How to Run Custom Scripts

- From the **Manage Telegraf Agents** tab, filter by **Agent Status › Agent Running**.
- Expand the drop-down arrow against the target machine on which the agent is installed. You see the **Custom Monitoring** section.
- Against the **Custom Script** option, click the vertical ellipsis and then click **Add**.
- From the **Manage Custom Script** dialog box, you can add and configure the Windows services to be monitored.

Instance Settings

| Option | Description |
|---|---|
| Status | Activate the custom script execution. |
| Display Name | Add a suitable name for the script. The `*` is an invalid character and must not be used in the name. |
| Filepath | Enter the path to the script file on the target machine.<br>Example:<br>For Linux machines:<br>`/opt/vmware/scripts/customscript.py`<br>For Windows machines:<br>`c:\scripts\customscript.ps1` |
| Prefix | Enter a prefix if necessary.<br>Example:<br>For Linux machines: python2.7, /bin/bash, or perl, etc<br>For Windows machines: *powershell -File* |
| Args | List the arguments in the script. |
| Timeout | Enter a script execution timeout on the VM. |

After you save the script, it appears under **Custom Script**. You can edit or delete scripts by clicking the **Edit** or **Delete** options from the vertical ellipsis against the custom script you added. After the scripts have been added and saved, click the drop down arrow against **Custom Script**, to view the list of scripts and their status.

**NOTE**

- The custom script must throw all errors in the format **`ERROR|<Error_message>`** for the error propagation to work. If the script does not throw an error in the given format, VMware Aria OperationsVMware Cloud Foundation Operations displays an error message **`Unable to parse the error message. Please check the endpoint`** in the user interface. This is by design, until cloud proxy propagates the exact error message.
- The bash script must start with `shebang (#!/bin/bash)`.

All Metrics Tab

When data is collected successfully, you can view the script as a metric for the target machine, in the **All Metrics** tab. The script metrics are created under an object called `Custom Script` which is a single object per target machine. All the metrics from the scripts for the target machine are placed under that `Custom Script` object that contains all the custom scripts you have created. You can view the output for the specific metric. The metric name under the `Scripts` folder is the display name that the user specifies while creating the script configuration. For example, if you set the display name as `Python script`, then a metric is created with the name **Python script** if data is collected successfully.

### Activate Remote Checks

You can activate remote checks such as ICMP Check, UDP Check, TCP Check, and HTTP Check.

- ICMP check sends a ping from the target machine to a specified resource.
- UDP check performs a UDP based connection check from the target machine to a specified resource.
- TCP check performs a connection check from the target machine to the port of a specified resource.
- HTTP check performs REST API calls from the target machine to a specified URL.

1. From the **Manage Telegraf Agents** tab, filter by **Agent Status › Agent Running**.
2. Expand the drop-down arrow against the relevant target machine on which the agent is installed. You see the **Custom Monitoring** section.
3. Against the **ICMP Check**, **UDP Check**, **TCP Check**, or **HTTP Check** options, click the vertical ellipsis and then click **Add**.
4. From the dialog box that appears on the right side, you can activate and configure the remote checks to be monitored.

   For configuration information, see Configuring Supported Application Services.

5. Click **Save**.

### Monitor Windows Services

After you install an agent on a target machine, you can monitor existing or custom Windows services that run on the target machine.

How to Monitor Windows Services

- From the **Manage Telegraf Agents** tab, filter by **Agent Status › Agent Running**.
- From the data grid, expand the drop-down arrow in front of the relevant Windows target machine name on which the agent is installed. You see the **Custom Monitoring** section.
- Against the **Services** option, click the vertical ellipsis and then click **Add**.
- From the **Manage Service Activation** dialog box, you can add and configure the Windows services to be monitored.

**Table 59: Instance Settings and Other Options**

| Option | Descriptions |
|---|---|
| Status | Activate the monitoring of the Windows service. |
| Display Name | Add a suitable name for the Windows service.<br>For new plugin activations, the VM name or the end point host name in the format <on *VM name*>, is automatically appended to the display name. For example, if the display name you enter is <system>, the VM name is automatically appended and the name is displayed as <system on *VM name*>.<br><br>If the display name was <system on abcd>, after an upgrade, <abcd> is replaced with the VM name or end point host name. |

*Table continued on next page*

*Continued from previous page*

| Option | Descriptions |
|---|---|
| | If the display name did not end with <on *text1*>, after an upgrade, <on *VM name*> is automatically appended to the existing display name.<br><br>The following are invalid characters and must not be used in the name: <, ", >, and \|. |
| Service Name | Enter a name of the Windows service you want to monitor. |

Save the settings to add the Windows service. To edit or delete Windows services, click the **Edit** or **Delete** options from the vertical ellipsis against Windows service you added. After the services have been added and saved, click the drop down arrow against **Services**, to view the list of Windows services and their status.

Metrics Tab

When data is collected successfully, you can view the metric from the **Manage Telegraf Agents** page, click the **Vertical Ellipsis › Go To Details › Metrics** tab. The metrics for the Windows service are created under an object called `Services` which is a single object per target machine.

### Monitor Linux Processes

After you install an agent on a target machine, you can monitor existing or custom Linux processes that run on the machine.

How to Monitor Linux Services

- From the **Manage Telegraf Agents** tab, filter by **Agent Status › Agent Running**.
- Expand the drop-down arrow against the relevant Linux machine on which the agent is installed. You see the **Custom Monitoring** section.
- Against the **Services** option, click the vertical ellipsis and then click **Add**.
- From the **Manage Service Activation** dialog box, you can add and configure the Linux services to be monitored.

**Table 60: Instance Settings and Other Options**

| Option | Description |
|---|---|
| Status | Activate or deactivate the monitoring of the Linux process. |
| Display Name | Add a suitable name for the Linux process you want to monitor.<br>For new plugin activations, the VM name or end point host name in the format <on *VM name*>, is automatically appended to the display name. For example, if the display name you enter is <system>, the VM name is automatically appended and the name is displayed as <system on *VM name*>.<br><br>If the display name was <system on abcd>, after an upgrade, <abcd> is replaced with the VM name or end point host name. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | If the display name did not end with <on *text1*>, after an upgrade, <on *VM name*> is automatically appended to the existing display name.<br><br>The following are invalid characters and must not be used in the name: <, ", >, and \|. |
| Filter Type | Select either **Executable Name**, **Regex Pattern**, or **Pid File** as the filter type from the drop-down menu. |
| Filter Value | The filter value could be a process executable name, a regex pattern, or a pid file absolute path. |

Save the settings to add the Linux service. To edit or delete Linux services, click the **Edit** or **Delete** options from the vertical ellipsis against Linux service you added. After the services have been added and saved, click the drop down arrow against **Services**, to view the list of Linux services and their status.

Metrics Tab

When data is collected successfully, you can view the metric from the **Manage Telegraf Agents** page, click the **Vertical Ellipsis › Go To Details › Metrics** tab. The metrics for the Linux process are created under an object called `Processes` which is a single object per target machine.

### Deactivate an Application Service

You can deactivate an application service to stop monitoring the application service that is sending data to VMware Aria OperationsVMware Cloud Foundation Operations.

**Prerequisite**

- If plugin deactivation requires the location of a file (for example, client certificates for SSL Trust) on the target machine, the location and the files should have appropriate read permissions for the *arcuser* to access those files.

  **NOTE**
  If the plugin displays a permission denied status, provide the *arcuser* with permissions to the file locations that you have specified during plugin activation.

### Deactivate an Application Service

To deactivate a plugin to stop monitoring the application service that is sending data to VMware Aria OperationsVMware Cloud Foundation Operations, complete the following steps:

1. From the left menu, click **Operations › Applications**. From the **Applications** panel, click **Manage Telegraf Agents**.
2. Filter by **Agent Status › Agent Running**.
3. Expand the drop-down arrow against the target machine on which the agent is installed. You see the **Services Discovered** section.
4. From the **Services Discovered** section, select a service that has been activated, click the vertical ellipsis and then click **Edit**.
5. Deactivate the application service from dialog box that is displayed on the right side.
6. Click **Save**.

For information on activating an application service, see Activate an Application Service.

## Summary of Discovered and Supported Operating Systems and Application Services

You can monitor application services and operating systems from VMware Aria OperationsVMware Cloud Foundation Operations and get insights into the services, processes and infrastructure. You can view a summary of discovered operating systems and services, supported operating systems, and supported services.

To monitor and view applications and operating systems, from the left menu, click **Operations › Configurations**, and then click the **Application Monitoring: Telegraf** tile from the right panel.

### Discovered Operating Systems and Services

You see the application services that are discovered on the end point virtual machines where the agents are installed. From the **Discovered Operating Systems and Services** section in the **Monitor Applications** page, click the <discovered> link under the name of the application service to perform lifecycle management actions for the agent and application services. For more information, see Additional Operations from the Manage Telegraf Agents Page .

### Supported Operating Systems

You see a list of supported operating systems for which VMware Aria OperationsVMware Cloud Foundation Operations collects metrics. For more information, see Supported Platforms.

### Supported Services

You see a list of supported services for which VMware Aria OperationsVMware Cloud Foundation Operations collects metrics. For more information, see Supported Application Services .

## Metrics Collected

Metrics are collected for operating systems, application services, remote checks, Linux processes, and Windows services.

## Operating System Metrics

Metrics are collected for Linux and Windows operating systems.

Linux Platforms

The following metrics are collected for Linux operating systems:

**Table 61: Metrics for Linux**

| Metric | Metric Category | KPI |
|---|---|---|
| <Instance name>\| Usage Idle | CPU | False |
| <Instance name>\| Usage IO-Wait | CPU | False |
| <Instance name>\|Time Active | CPU | True |
| <Instance name>\|Time Guest | CPU | False |
| <Instance name>\|Time Guest Nice | CPU | False |
| <Instance name>\|Time Idle | CPU | False |
| <Instance name>\|Time IO-Wait | CPU | False |
| <Instance name>\|Time IRQ | CPU | True |

*Table continued on next page*

*Continued from previous page*

| Metric | Metric Category | KPI |
|---|---|---|
| <Instance name>\|Time Nice | CPU | False |
| <Instance name>\|Time Soft IRQ | CPU | True |
| <Instance name>\|Time Steal | CPU | False |
| <Instance name>\|Time System | CPU | False |
| <Instance name>\|Time User | CPU | True |
| <Instance name>\|Usage Active (%) | CPU | True |
| <Instance name>\|Usage Guest (%) | CPU | False |
| <Instance name>\|Usage Guest Nice (%) | CPU | False |
| <Instance name>\|Usage IRQ (%) | CPU | True |
| <Instance name>\|Usage Nice (%) | CPU | False |
| <Instance name>\|Usage Soft IRQ (%) | CPU | True |
| <Instance name>\|Usage Steal (%) | CPU | False |
| <Instance name>\|Usage System (%) | CPU | True |
| <Instance name>\|Usage User (%) | CPU | True |
| CPU Load1 (%) | CPU Load | False |
| CPU Load15 (%) | CPU Load | False |
| CPU Load5 (%) | CPU Load | False |
| <Instance name>\|IO Time | Disk IO | False |
| <Instance name>\|Read Time | Disk IO | False |
| <Instance name>\|Reads | Disk IO | False |
| <Instance name>\|Write Time | Disk IO | False |
| <Instance name>\|Writes | Disk IO | False |
| <Instance name>\|Disk Free | Disk | False |
| <Instance name>\|Disk Total | Disk | False |
| <Instance name>\|Disk Used (%) | Disk | False |
| Cached | Memory | False |
| Free | Memory | False |
| Inactive | Memory | False |
| Total | Memory | True |
| Used | Memory | True |
| Used Percent | Memory | True |
| Blocked | Processes | True |
| Dead | Processes | False |
| Running | Processes | False |
| Sleeping | Processes | False |
| Stopped | Processes | False |
| Zombies | Processes | False |
| Free | Swap | False |
| In | Swap | False |
| Out | Swap | False |
| Total | Swap | True |

*Table continued on next page*

*Continued from previous page*

| Metric | Metric Category | KPI |
|---|---|---|
| Used | Swap | True |
| Used Percent | Swap | True |
| Telegraf Availability | None | False |

Windows Platforms

The following metrics are collected for Windows operating systems:

**Table 62: Metrics for Windows**

| Metric | Metric Category | KPI |
|---|---|---|
| Idle Time | CPU | False |
| Interrupt Time | CPU | False |
| Interrupts persec | CPU | True |
| Privileged Time | CPU | False |
| Processor Time | CPU | False |
| User Time | CPU | False |
| DPC Time (%) | CPU | False |
| Usage Guest (%) | CPU | False |
| Usage System (%) | CPU | False |
| Usage User (%) | CPU | False |
| Avg. Disk Bytes Read | Disk | False |
| Avg. Disk sec Read | Disk | False |
| Avg. Disk sec Write | Disk | False |
| Avg. Disk Write Queue Length | Disk | False |
| Avg. Disk Read Queue Length | Disk | False |
| Disk Read Time | Disk | False |
| Disk Write Time | Disk | False |
| Free Megabytes | Disk | False |
| Free Space | Disk | False |
| Idle Time | Disk | False |
| Split IO persec | Disk | False |
| Available Bytes | Memory | True |
| Cache Bytes | Memory | False |
| Cache Faults persec | Memory | False |
| Committed Bytes | Memory | True |
| Demand Zero Faults persec | Memory | False |
| Page Faults persec | Memory | True |
| Pages persec | Memory | False |
| Pool Nonpaged Bytes | Memory | True |
| Pool Paged Bytes | Memory | False |
| Transition Faults persec | Memory | False |
| Total (bytes) | Memory | False |
| Used (bytes) | Memory | False |

*Table continued on next page*

*Continued from previous page*

| Metric | Metric Category | KPI |
|---|---|---|
| Used Percent(%) | Memory | False |
| Bytes Received persec | Network | False |
| Bytes Sent persec | Network | False |
| Packets Outbound Discarded | Network | False |
| Packets Outbound Errors | Network | False |
| Packets Received Discarded | Network | False |
| Packets Received Errors | Network | False |
| Packets Received persec | Network | False |
| Packets Sent persec | Network | False |
| Elapsed Time | Process | False |
| Handle Count | Process | False |
| IO Read Bytes persec | Process | False |
| IO Read Operations persec | Process | False |
| IO Write Bytes persec | Process | False |
| IO Write Operations persec | Process | False |
| Privileged Time | Process | False |
| Processor Time | Process | False |
| Thread Count | Process | False |
| User Time | Process | False |
| Context Switches persec | System | False |
| Processes | System | False |
| Processor Queue Length | System | False |
| System Calls persec | System | False |
| System Up Time | System | False |
| Threads | System | False |
| Used Percent (%) | Swap | False |
| Total (bytes) | Swap | False |
| Telegraf Availability | None | False |

**Application Service Metrics**

Metrics are collected for 23 application services.

## Active Directory Metrics

Metrics are collected for the Active Directory application service.

**Table 63: Active Directory Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Database Cache % Hit (%) | Active Directory Database | True |
| Database Cache Page Faults/sec | Active Directory Database | True |
| Database Cache Size | Active Directory Database | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Data Lookups | Active Directory DFS Replication | False |
| Database Commits | Active Directory DFS Replication | True |
| Avg Response Time | Active Directory DFSN | True |
| Requests Failed | Active Directory DFSN | False |
| Requests Processed | Active Directory DFSN | False |
| Dynamic Update Received | Active Directory DNS | False |
| Dynamic Update Rejected | Active Directory DNS | False |
| Recursive Queries | Active Directory DNS | False |
| Recursive Queries Failure | Active Directory DNS | False |
| Secure Update Failure | Active Directory DNS | False |
| Total Query Received | Active Directory DNS | True |
| Total Response Sent | Active Directory DNS | True |
| Digest Authentications | Active Directory Security System-Wide Statistics | True |
| Kerberos Authentications | Active Directory Security System-Wide Statistics | True |
| NTLM Authentications | Active Directory Security System-Wide Statistics | True |
| Directory Services:<InstanceName>\| Base Searches persec | Active Directory Services | False |
| Directory Services:<InstanceName>\| Database adds persec | Active Directory Services | False |
| Directory Services:<InstanceName>\| Database deletes persec | Active Directory Services | False |
| Directory Services<InstanceName>\| Database modifys/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| Database recycles/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Inbound Bytes Total/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Inbound Objects/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Outbound Bytes Total/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Outbound Objects/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Pending Replication Operations | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Pending Replication Synchronizations | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Sync Requests Made | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Sync Requests Successful | Active Directory Services | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Directory Services<InstanceName>\| DS Client Binds/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| DS Directory Reads/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DS Directory Searches/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| DS Server Binds/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| DS Threads in Use | Active Directory Services | True |
| Directory Services:<InstanceName>\| LDAP Active Threads | Active Directory Services | False |
| Directory Services:<InstanceName>\| LDAP Client Sessions | Active Directory Services | True |
| Directory Services<InstanceName>\| LDAP Closed Connections/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| LDAP New Connections/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| LDAP Searches/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| LDAP Successful Binds/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| LDAP UDP operations/sec | Active Directory Services | False |
| Directory Services:<InstanceName>\| LDAP Writes/sec | Active Directory Services | False |
| Application Availability | Active Directory | False |

## Apache HTTPD

Metrics are collected for the Apache HTTPD application service.

**Table 64: Apache Tomcat**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Tomcat Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Tomcat Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Tomcat Server | False |
| Class Loading\|Loaded Class Count | Tomcat Server | False |
| Class Loading\|Total Loaded Class Count | Tomcat Server | False |
| Class Loading\|Unloaded Class Count | Tomcat Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Tomcat Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| File Descriptor Usage\|Open File Descriptor Count | Tomcat Server | False |
| Garbage Collection:<InstanceName>\|Total Collection Count | Tomcat Server | False |
| Garbage Collection:<InstanceName>\|Total Collection Time | Tomcat Server | True |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Tomcat Server | False |
| JVM Memory\|Number of Object Pending Finalization Count | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Used Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Used Memory | Tomcat Server | False |
| Process CPU Usage (%) | Tomcat Server | True |
| System CPU Usage (%) | Tomcat Server | True |
| System Load Average (%) | Tomcat Server | True |
| Threading\|Thread Count | Tomcat Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Uptime | Tomcat Server | True |
| Application Availability | Tomcat Server | False |
| JSP Count | Tomcat Server Web Module | False |
| JSP Reload Count | Tomcat Server Web Module | False |
| JSP Unload Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Error Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Processing Time | Tomcat Server Web Module | False |
| Cache : Hit Count | Tomcat Server Web Module | False |
| Cache : Lookup Count | Tomcat Server Web Module | False |
| Current Thread Count | Tomcat Server Global Request Processor | True |
| Current Threads Busy | Tomcat Server Global Request Processor | True |
| errorRate | Tomcat Server Global Request Processor | False |
| Total Request Bytes Received | Tomcat Server Global Request Processor | False |
| Total Request Bytes Sent | Tomcat Server Global Request Processor | False |
| Total Request Count | Tomcat Server Global Request Processor | True |
| Total Request Error Count | Tomcat Server Global Request Processor | True |
| Total Request Processing Time | Tomcat Server Global Request Processor | False |

## Microsoft SQL Server Metrics

Metrics are collected for the Microsoft SQL Server application service.

**Table 65: MS SQL Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| CPU<InstanceName>\|CPU Usage (%) | Microsoft SQL Server | False |
| Database IO\|Rows Reads Bytes/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Reads/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Writes Bytes/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Writes/Sec | Microsoft SQL Server | False |
| Performance\|Access Methods\|Full Scans per second | Microsoft SQL Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Performance\|Access Methods\|Index Searches | Microsoft SQL Server | False |
| Performance\|Access Methods\|Page Splits per second | Microsoft SQL Server | False |
| Performance\|Broker Activation\|Stored Procedures Invoked per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Buffer cache hit ratio (%) | Microsoft SQL Server | True |
| Performance\|Buffer Manager\| Checkpoint Pages/sec | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Lazy writes per second | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Page life expectancy | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Page lookups per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Page reads per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Page writes per second | Microsoft SQL Server | False |
| Performance\|Databases\|Active Transactions | Microsoft SQL Server | True |
| Performance\|Databases\|Data File(s) Size | Microsoft SQL Server | True |
| Performance\|Databases\|Log Bytes Flushed/Sec | Microsoft SQL Server | False |
| Performance\|Databases\|Log File(s) Size | Microsoft SQL Server | False |
| Performance\|Databases\|Log File(s) Used Size | Microsoft SQL Server | False |
| Performance\|Databases\|Log Flush Wait Time | Microsoft SQL Server | False |
| Performance\|Databases\|Log Flushes per second | Microsoft SQL Server | False |
| Performance\|Databases\| Transactions per second | Microsoft SQL Server | False |
| Performance\|Databases\|Write Transactions per second | Microsoft SQL Server | False |
| Performance\|Databases\|XTP Memory Used | Microsoft SQL Server | False |
| Performance\|General Statistics\|Active temp Tables | Microsoft SQL Server | False |
| Performance\|General Statistics\| Logins per second | Microsoft SQL Server | False |
| Performance\|General Statistics\| Logouts per second | Microsoft SQL Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Performance\|General Statistics\| Processes Blocked | Microsoft SQL Server | False |
| Performance\|General Statistics\|Temp Tables Creation Rate | Microsoft SQL Server | False |
| Performance\|General Statistics\|User Connections | Microsoft SQL Server | False |
| Performance\|Locks\|Average Wait Time | Microsoft SQL Server | False |
| Performance\|Locks\|Lock Requests per second | Microsoft SQL Server | False |
| Performance\|Locks\|Lock Wait Time | Microsoft SQL Server | True |
| Performance\|Locks\|Lock Waits per second | Microsoft SQL Server | True |
| Performance\|Locks\|Number of Deadlocks per second | Microsoft SQL Server | True |
| Performance\|Memory Manager\| Connection Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\|Lock Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\|Log Pool Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\| Memory Grants Pending | Microsoft SQL Server | True |
| Performance\|Memory Manager\|SQL Cache Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\| Target Server Memory | Microsoft SQL Server | True |
| Performance\|Memory Manager\|Total Server Memory | Microsoft SQL Server | True |
| Performance\|Resource Pool Stats\| internal\|Active memory grant amount | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|CPU Usage Percentage (%) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read Bytes per second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read IO | Microsoft SQL Server | False |
| Wait Stats:<InstanceName>\|Wait Time (ms) | Microsoft SQL Server | False |
| Wait Stats<InstanceName>\|Number of Waiting tasks (ms) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read IO Throttled Per Second | Microsoft SQL Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Performance\|Resource Pool Stats\| internal\|Disk Write Bytes per second (Bps) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Write IO Throttled per second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Used Memory | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| Batch Requests Per Second | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| SQL Compilations per second | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| SQL Re-Compilations per second | Microsoft SQL Server | False |
| Performance\|Transactions \| Free space in tempdb (KB) | Microsoft SQL Server | False |
| Performance\|Transactions \| Transactions | Microsoft SQL Server | False |
| Performance\|Transactions \| Version Store Size (KB) | Microsoft SQL Server | False |
| Performance\|User Settable Counter \| User Counter 0 to 10 | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Active Requests | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Blocked Tasks | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|CpU Usage (%) | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Queued Requests | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Request Completed/sec | Microsoft SQL Server | False |
| Application Availability | Microsoft SQL Server | False |

There are no metrics collected for Microsoft SQL Server Database.

## PostgreSQL

Metrics are collected for the PostgreSQL application service.

**Table 66: PostgreSQL**

| Metric Name | Category | KPI |
|---|---|---|
| Buffers\|Buffers Allocated | PostgreSQL | False |
| Buffers\|Buffers Written by Backend | PostgreSQL | True |
| Buffers\|Buffers Written by Background Writer | PostgreSQL | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Buffers\|Buffers Written During Checkpoints | PostgreSQL | True |
| Buffers\|fsync Call Executed by Backend | PostgreSQL | False |
| Checkpoints\|Checkpoints sync time | PostgreSQL | False |
| Checkpoints\|Checkpoints write time | PostgreSQL | False |
| Checkpoints\|Requested checkpoints performed count | PostgreSQL | False |
| Checkpoints\|Scheduled checkpoints performed count | PostgreSQL | False |
| Clean scan stopped count | PostgreSQL | False |
| Application Availability | PostgreSQL | False |
| Disk Blocks\|Blocks Cache Hits | PostgreSQL Database | False |
| Disk Blocks\|Blocks Read | PostgreSQL Database | False |
| Disk Blocks\|Blocks Read Time | PostgreSQL Database | False |
| Disk Blocks\|Blocks Write Time | PostgreSQL Database | False |
| Statistics\|Backends Connected | PostgreSQL Database | False |
| Statistics\|Data Written by Queries | PostgreSQL Database | True |
| Statistics\|Deadlocks Detected | PostgreSQL Database | True |
| Statistics\|Queries Cancelled | PostgreSQL Database | True |
| Statistics\|Temp Files Created by Queries | PostgreSQL Database | False |
| Transactions\|Transactions Committed | PostgreSQL Database | True |
| Transactions\|Transactions Rolled Back | PostgreSQL Database | True |
| Tuples\|Tuples Deleted | PostgreSQL Database | True |
| Tuples\|Tuples Fetched | PostgreSQL Database | True |
| Tuples\|Tuples Inserted | PostgreSQL Database | True |
| Tuples\|Tuples Returned | PostgreSQL Database | True |
| Tuples\|Tuples Updated | PostgreSQL Database | True |

## Microsoft IIS Metrics

Metrics are collected for the Microsoft IIS application service.

**Table 67: IIS Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Service Request Queues<InstanceName>AppPool\|CurrentQueueSize | IIS HTTP Service Request Queues | True |
| HTTP Service Request Queues<InstanceName>AppPool\|RejectedRequests | IIS HTTP Service Request Queues | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Web Services<InstanceName> Web Site\|Bytes Received | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Bytes Sent/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Bytes Total/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Connection Attempts/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Current Connections | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Get Requests/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Locked Errors/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Not Found Errors/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Post Requests/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Service Uptime | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Total Bytes Sent | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Total Get Requests | IIS Web Services | True |
| Web Services<InstanceName> Web Site\|Total Post Requests | IIS Web Services | True |
| Web Services<InstanceName> Web Site\|Total Put Requests | IIS Web Services | False |
| Current File Cache Memory Usage (bytes) | IIS Web Services Cache | False |
| File Cache Hits Percent (%) | IIS Web Services Cache | False |
| Kernel URI Cache Hits Percent (%) | IIS Web Services Cache | False |
| Kernel URI Cache Misses | IIS Web Services Cache | False |
| Total Flushed URIs | IIS Web Services Cache | False |
| URI Cache Hits | IIS Web Services Cache | False |
| URI Cache Hits Percent (%) | IIS Web Services Cache | False |
| URI Cache Misses | IIS Web Services Cache | False |
| ASP.NET<InstanceName>\|Application Restarts | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Request Wait Time | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Requests Current | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Requests Queued | IIS ASP.NET | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| ASP.NET<InstanceName>\|Requests Rejected | IIS ASP.NET | True |
| MS.NET<InstanceName>\|Allocated Bytes/sec | MS.NET | True |
| MS.NET<InstanceName>\|Current Queue Length | MS.NET | False |
| MS.NET<InstanceName>\|Finalization Survivors | MS.NET | False |
| MS.NET<InstanceName>\|Gen 0 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 0 heap size | MS.NET | False |
| MS.NET<InstanceName>\|Gen 1 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 1 heap size | MS.NET | False |
| MS.NET<InstanceName>\|Gen 2 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 2 heap size | MS.NET | False |
| MS.NET<InstanceName>\|IL Bytes Jitted / sec | MS.NET | False |
| MS.NET<InstanceName>\|Induced GC | MS.NET | False |
| MS.NET<InstanceName>\|Large Object Heap size | MS.NET | False |
| MS.NET<InstanceName>\|No of current logical Threads | MS.NET | True |
| MS.NET<InstanceName>\|No of current physical Threads | MS.NET | True |
| MS.NET<InstanceName>\|No of current recognized threads | MS.NET | False |
| MS.NET<InstanceName>\|No of Exceps Thrown / sec | MS.NET | True |
| MS.NET<InstanceName>\|No of total recognized threads | MS.NET | False |
| MS.NET<InstanceName>\|Percent Time in Jit | MS.NET | False |
| MS.NET<InstanceName>\|Pinned Objects | MS.NET | False |
| MS.NET<InstanceName>\|Stack Walk Depth | MS.NET | False |
| MS.NET<InstanceName>\|Time in RT checks | MS.NET | False |
| MS.NET<InstanceName>\|Time Loading | MS.NET | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| MS.NET<InstanceName>\|Total No of Contentions | MS.NET | False |
| MS.NET<InstanceName>\|Total Runtime Checks | MS.NET | True |
| Application Availability | Microsoft IIS | False |

## MS Exchange Metrics

Metrics are collected for the MS Exchange application service.

**Table 68: MS Exchange Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Active Manager Server\|Active Manager Role | MS Exchange | False |
| Active Manager Server\|Database State Info Writes per second | MS Exchange | False |
| Active Manager Server\| GetServerForDatabase Server-Side Calls | MS Exchange | False |
| Active Manager Server\|Server-Side Calls per second | MS Exchange | True |
| Active Manager Server\|Total Number of Databases | MS Exchange | True |
| ActiveSync\|Average Request Time | MS Exchange | True |
| ActiveSync\|Current Requests | MS Exchange | False |
| ActiveSync\|Mailbox Search Total | MS Exchange | False |
| ActiveSync\|Ping Commands Pending | MS Exchange | False |
| ActiveSync\|Requests per second | MS Exchange | True |
| ActiveSync\|Sync Commands per second | MS Exchange | True |
| ASP.NET\|Application Restarts | MS Exchange | False |
| ASP.NET\|Request Wait Time | MS Exchange | True |
| ASP.NET\|Worker Process Restarts | MS Exchange | False |
| Autodiscover Service\|Requests per second | MS Exchange | True |
| Availability Service\|Average Time to Process a Free Busy Request | MS Exchange | True |
| Outlook Web Access\|Average Search Time | MS Exchange | True |
| Outlook Web Access\|Requests per second | MS Exchange | False |
| Outlook Web Access\|Current Unique Users | MS Exchange | False |
| Application Availability | MS Exchange | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Performance\|Database Cache Hit (%) | MS Exchange Database | False |
| Performance\|Database Page Fault Stalls per second | MS Exchange Database | True |
| Performance\|I/O Database Reads Average Latency | MS Exchange Database | True |
| Performance\|I/O Database Writes Average Latency | MS Exchange Database | True |
| Performance\|I/O Log Reads Average Latency | MS Exchange Database | False |
| Performance\|I/O Log Writes Average Latency | MS Exchange Database | False |
| Performance\|Log Record Stalls per second | MS Exchange Database | False |
| Performance\|Log Threads Waiting | MS Exchange Database | False |
| Performance\|I/O Database Reads Average Latency | MS Exchange Database Instance | False |
| Performance\|I/O Database Writes Average Latency | MS Exchange Database Instance | False |
| Performance\|Log Record Stalls per second | MS Exchange Database Instance | False |
| Performance\|Log Threads Waiting | MS Exchange Database Instance | False |
| Performance\|LDAP Read Time | MS Exchange Domain Controller | False |
| Performance\|LDAP Search Time | MS Exchange Domain Controller | False |
| Performance\|LDAP Searches Timed Out per minute | MS Exchange Domain Controller | False |
| Performance\|Long Running LDAP Operations per minute | MS Exchange Domain Controller | False |
| Performance\|Connection Attempts per second | MS Exchange Web Server | True |
| Performance\|Current Connections | MS Exchange Web Server | False |
| Performance\|Other Request Methods per second | MS Exchange Web Server | False |
| Process\|Handle Count | MS Exchange Windows Service | False |
| Process\|Memory Allocated | MS Exchange Windows Service | False |
| Process\|Processor Time (%) | MS Exchange Windows Service | True |
| Process\|Thread Count | MS Exchange Windows Service | False |
| Process\|Virtual Memory Used | MS Exchange Windows Service | False |
| Process\|Working Set | MS Exchange Windows Service | False |

## JBoss Server Metrics

Metrics are collected for the JBoss Server application service.

**Table 69: JBoss Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Jboss Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Jboss Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Jboss Server | False |
| Class Loading\|Loaded Class Count | Jboss Server | False |
| Class Loading\|Total Loaded Class Count | Jboss Server | False |
| Class Loading\|Unloaded Class Count | Jboss Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Jboss Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Jboss Server | False |
| Http Listener<InstanceName>\|Bytes Received | Jboss Server | False |
| Http Listener<InstanceName>\|Bytes Sent | Jboss Server | False |
| Http Listener<InstanceName>\|Error Count | Jboss Server | False |
| Http Listener<InstanceName>\| Request Count | Jboss Server | False |
| Https Listener<InstanceName>\|Bytes Received | Jboss Server | False |
| Https Listener<InstanceName>\|Bytes Sent | Jboss Server | False |
| Https Listener<InstanceName>\|Error Count | Jboss Server | False |
| Https Listener<InstanceName>\| Request Count | Jboss Server | False |
| Process CPU Usage (%) | Jboss Server | False |
| System CPU Usage (%) | Jboss Server | False |
| System Load Average (%) | Jboss Server | False |
| Threading\|Daemon Thread Count | Jboss Server | False |
| Threading\|Peak Thread Count | Jboss Server | False |
| Threading\|Thread Count | Jboss Server | False |
| Threading\|Total Started Thread Count | Jboss Server | False |
| Uptime | Jboss Server | False |
| UTILIZATION\|Heap Memory Usage | Jboss Server | False |
| Application Availability | Jboss Server | False |
| Garbage Collection<InstanceName>\| Total Collection Count | Jboss JVM Garbage Collector | False |
| Garbage Collection<InstanceName>\| Total Collection Time | Jboss JVM Garbage Collector | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Heap Memory Usage\| Committed Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Jboss JVM Memory | True |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Jboss JVM Memory | False |
| JVM Memory\|Object Pending Finalization Count | Jboss JVM Memory | True |
| UTILIZATION\|Active Count | Jboss Datasource Pool | False |
| UTILIZATION\|Available Count | Jboss Datasource Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Used Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Peak Usage\|Used Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\| Usage\|Used Memory | Jboss JVM Memory Pool | False |

## RabbitMQ Metrics

Metrics are collected for the RabbitMQ application service.

**Table 70: RabbitMQ Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| CPU\|Limit | RabbitMQ | False |
| CPU\|Used | RabbitMQ | True |
| Disk\|Free | RabbitMQ | False |
| Disk\|Free limit | RabbitMQ | False |
| FileDescriptor\|Total | RabbitMQ | False |
| FileDescriptor\|Used | RabbitMQ | False |
| Memory\|Limit | RabbitMQ | False |
| Memory\|Used | RabbitMQ | True |
| Messages\|Acked | RabbitMQ | False |
| Messages\|Delivered | RabbitMQ | False |
| Messages\|Delivered get | RabbitMQ | False |
| Messages\|Published | RabbitMQ | False |
| Messages\|Ready | RabbitMQ | False |
| Messages\|Unacked | RabbitMQ | False |
| Socket\|Limit | RabbitMQ | False |
| Socket\|Used | RabbitMQ | True |
| UTILIZATION\|Channels | RabbitMQ | True |
| UTILIZATION\|Connections | RabbitMQ | True |
| UTILIZATION\|Consumers | RabbitMQ | True |
| UTILIZATION\|Exchanges | RabbitMQ | True |
| UTILIZATION\|Messages | RabbitMQ | True |
| UTILIZATION\|Queues | RabbitMQ | True |
| Application Availability | RabbitMQ | False |
| Messages\|Publish in | RabbitMQ Exchange | False |
| Messages\|Publish out | RabbitMQ Exchange | False |
| Consumer Utilisation | RabbitMQ Queue | False |
| Consumers | RabbitMQ Queue | False |
| Memory | RabbitMQ Queue | False |
| Messages\|Ack | RabbitMQ Queue | False |
| Messages\|Ack rate | RabbitMQ Queue | False |
| Messages\|Deliver | RabbitMQ Queue | False |
| Messages\|Deliver get | RabbitMQ Queue | False |
| Messages\|Persist | RabbitMQ Queue | False |
| Messages\|Publish | RabbitMQ Queue | False |
| Messages\|Publish rate | RabbitMQ Queue | False |
| Messages\|Ram | RabbitMQ Queue | False |
| Messages\|Ready | RabbitMQ Queue | False |
| Messages\|Redeliver | RabbitMQ Queue | False |
| Messages\|Redeliver rate | RabbitMQ Queue | False |
| Messages\|Space | RabbitMQ Queue | False |
| Messages\|Unack | RabbitMQ Queue | False |
| Messages\|Unacked | RabbitMQ Queue | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Messages | RabbitMQ Queue | False |

There are no metrics collected for RabbitMQ Virtual Host.

## MySQL Metrics

Metrics are collected for the MySQL application service.

**Table 71: MySQL Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Aborted connection count | MySQL | True |
| Connection count | MySQL | True |
| Event wait average time | MySQL | False |
| Event wait count | MySQL | False |
| Binary Files\|Binary Files Count | MySQL | False |
| Binary Files\|Binary Size Bytes | MySQL | False |
| Global Status\|Aborted Clients | MySQL | False |
| Global Status\|Binlog Cache Disk Use | MySQL | False |
| Global Status\|Bytes Received | MySQL | False |
| Global Status\|Bytes Sent | MySQL | False |
| Global Status\|Connection Errors Accept | MySQL | False |
| Global Status\|Connection Errors Internal | MySQL | False |
| Global Status\|Connection Errors Max Connections | MySQL | False |
| Global Status\|Queries | MySQL | False |
| Global Status\|Threads Cached | MySQL | False |
| Global Status\|Threads Connected | MySQL | False |
| Global Status\|Threads Running | MySQL | False |
| Global Status\|Uptime | MySQL | False |
| Global Variables\|Delayed Insert Limit | MySQL | False |
| Global Variables\|Delayed Insert Timeout | MySQL | False |
| Global Variables\|Delayed Queue Size | MySQL | False |
| Global Variables\|Max Connect Errors | MySQL | False |
| Global Variables\|Max Connections | MySQL | False |
| Global Variables\|Max Delayed Threads | MySQL | False |
| Global Variables\|Max Error Count | MySQL | False |
| InnoDB\|All deadlock count | MySQL | False |
| InnoDB\|Buffer Pool Bytes Data | MySQL | False |
| InnoDB\|Buffer Pool Bytes Data | MySQL | False |
| InnoDB\|Buffer Pool Bytes Dirty | MySQL | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| InnoDB\|Buffer Pool Dump Status | MySQL | False |
| InnoDB\|Buffer Pool Load Status | MySQL | False |
| InnoDB\|Buffer Pool Pages Data | MySQL | False |
| InnoDB\|Buffer Pool Pages Dirty | MySQL | False |
| InnoDB\|Buffer Pool Pages Flushed | MySQL | False |
| InnoDB\|Buffer pool size | MySQL | True |
| InnoDB\|Checksums | MySQL | False |
| InnoDB\|Open file count | MySQL | False |
| InnoDB\|Row lock average time | MySQL | False |
| InnoDB\|Row lock current waits | MySQL | False |
| InnoDB\|Row lock maximum time | MySQL | False |
| InnoDB\|Row lock time | MySQL | False |
| InnoDB\|Row lock waits | MySQL | True |
| InnoDB\|Table lock count | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Delete | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Fetch | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Insert | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Update | MySQL | False |
| Process List\|Connections | MySQL | False |
| Application Availability | MySQL | False |
| IO waits average time | MySQL Database | False |
| IO waits count | MySQL Database | True |
| Read high priority average time | MySQL Database | False |
| Read high priority count | MySQL Database | False |
| Write concurrent insert average time | MySQL Database | False |
| Write concurrent insert count | MySQL Database | False |

## NGINX Metrics

Metrics are collected for the NGINX application service.

**Table 72: NGINX Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Status Info\|Accepts | Nginx | True |
| HTTP Status Info\|Active connections | Nginx | False |
| HTTP Status Info\|Handled | Nginx | True |
| HTTP Status Info\|Reading | Nginx | False |
| HTTP Status Info\|Requests | Nginx | False |
| HTTP Status Info\|Waiting | Nginx | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Status Info\|Writing | Nginx | False |
| Application Availability | Nginx | False |

## SharePoint Metrics

Metrics are collected for the SharePoint Server application service.

**Table 73: SharePoint Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Sharepoint Foundation\|Active Threads | SharePoint Server | True |
| Sharepoint Foundation\|Current Page Requests | SharePoint Server | False |
| Sharepoint Foundation\|Executing SQL Queries | SharePoint Server | False |
| Sharepoint Foundation\|Executing Time/Page Request | SharePoint Server | True |
| Sharepoint Foundation\|Incoming Page Requests Rate | SharePoint Server | False |
| Sharepoint Foundation\|Object Cache Hit Count | SharePoint Server | False |
| Sharepoint Foundation\|Reject Page Requests Rate | SharePoint Server | False |
| Sharepoint Foundation\|Responded Page Requests Rate | SharePoint Server | True |
| SQL query executing time | SharePoint Server | False |
| Application Availability | SharePoint Server | False |
| Network\|Received Data Rate | SharePoint Web Server | True |
| Network\|Sent Data Rate | SharePoint Web Server | True |
| Process\|Processor Time (%) | SharePoint Windows Service | False |
| Process\|Threads | SharePoint Windows Service | False |

## Oracle WebLogic Server Metrics

Metrics are collected for the Oracle WebLogic Server application service.

**Table 74: Oracle WebLogic Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Process Cpu Load | Oracle WebLogic Server | True |
| UTILIZATION\|System Cpu Load | Oracle WebLogic Server | False |
| UTILIZATION\|System Load Average | Oracle WebLogic Server | False |
| Application Availability | Oracle WebLogic Server | False |
| UTILIZATION\|Collection Time | Weblogic Garbage Collector | True |
| UTILIZATION\|Connections HighCount | Weblogic JMS Runtime | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|JMS Servers TotalCount | Weblogic JMS Runtime | False |
| UTILIZATION\|Active Total Count Used | Weblogic JTA Runtime | False |
| UTILIZATION\|Active Transactions TotalCount | Weblogic JTA Runtime | False |
| UTILIZATION\|Transaction Abandoned TotalCount | Weblogic JTA Runtime | True |
| UTILIZATION\|Transaction RolledBack App TotalCount | Weblogic JTA Runtime | True |
| UTILIZATION\|Heap Memory Usage | Weblogic JVM Memory | True |
| UTILIZATION\|Non Heap Memory Usage | Weblogic JVM Memory | False |
| UTILIZATION\|Peak Usage | Weblogic JVM Memory Pool | True |
| UTILIZATION\|Usage | Weblogic JVM Memory Pool | False |
| UTILIZATION\|UpTime | Weblogic JVM Runtime | False |

## Pivotal TC Server Metrics

Metrics are collected for the Pivotal TC Server application service.

**Table 75: Pivotal TC Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Pivotal TC Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Pivotal TC Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Pivotal TC Server | False |
| Class Loading\|Loaded Class Count | Pivotal TC Server | False |
| Class Loading\|Total Loaded Class Count | Pivotal TC Server | False |
| Class Loading\|Unloaded Class Count | Pivotal TC Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Pivotal TC Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Pivotal TC Server | False |
| Garbage Collection:<InstanceName>\| Total Collection Count | Pivotal TC Server | False |
| Garbage Collection:<InstanceName>\|Total Collection Time | Pivotal TC Server | False |
| Process CPU Usage (%) | Pivotal TC Server | True |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Pivotal TC Server | True |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Pivotal TC Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Pivotal TC Server | True |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Pivotal TC Server | True |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Pivotal TC Server | True |
| JVM Memory\|Number of Object Pending Finalization Count | Pivotal TC Server | True |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Committed Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Used Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Committed Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Used Memory | Pivotal TC Server | False |
| Process CPU Usage (%) | Pivotal TC Server | True |
| System CPU Usage (%) | Pivotal TC Server | True |
| Uptime | Pivotal TC Server | True |
| Threading\|Thread Count | Pivotal TC Server | False |
| System Load Average | Pivotal TC Server | False |
| Application Availability | Pivotal TC Server | False |
| Current Thread Count | Pivotal TC Server Thread Pool | False |
| Current Threads Busy | Pivotal TC Server Thread Pool | True |
| Total Request Bytes Received | Pivotal TC Server Thread Pool | False |
| Total Request Bytes Sent | Pivotal TC Server Thread Pool | False |
| Total Request Count | Pivotal TC Server Thread Pool | True |
| Total Request Error Count | Pivotal TC Server Thread Pool | True |
| Total Request Processing Time | Pivotal TC Server Thread Pool | True |
| JSP Count | Pivotal TC Server Web Module | False |
| JSP Reload Count | Pivotal TC Server Web Module | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JSP Unload Count | Pivotal TC Server Web Module | False |

## ActiveMQ Metrics

Metrics are collected for the ActiveMQ application service.

**Table 76: ActiveMQ Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Active MQ | False |
| Buffer Pool<InstanceName>\| Memory Used | Active MQ | False |
| Buffer Pool<InstanceName>\|Total Capacity | Active MQ | False |
| Class Loading\|Loaded Class Count | Active MQ | False |
| Class Loading\|Unloaded Class Count | Active MQ | False |
| Class Loading\|Total Loaded Class Count | Active MQ | False |
| File Descriptor Usage\|Max File Descriptor Count | Active MQ | False |
| File Descriptor Usage\|Open File Descriptor Count | Active MQ | False |
| Garbage Collection<InstanceName>\| Total Collection Count | Active MQ | False |
| Garbage Collection<InstanceName>\| Total Collection Time | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Committed Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Initial Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Maximum Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Used Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Committed Memory | Active MQ | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory Pool<InstanceName>\|Usage\|Initial Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Usage\|Maximum Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Usage\|Used Memory | Active MQ | False |
| Application Availability | Active MQ | False |
| Threading\|Thread Count | Active MQ | False |
| Uptime | Active MQ | False |
| UTILIZATION\|Process CpuLoad | Active MQ | False |
| UTILIZATION\|Memory Limit | ActiveMQ Broker | True |
| UTILIZATION\|Memory Percent Usage (%) | ActiveMQ Broker | True |
| UTILIZATION\|Store Limit | ActiveMQ Broker | False |
| UTILIZATION\|Store Percent Usage (%) | ActiveMQ Broker | False |
| UTILIZATION\|Temp Limit | ActiveMQ Broker | False |
| UTILIZATION\|Temp Percent Usage (%) | ActiveMQ Broker | False |
| UTILIZATION\|Total Consumer Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Dequeue Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Enqueue Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Message Count | ActiveMQ Broker | True |
| JVM Memory\|Heap Memory Usage\|Initial Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Maximum Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Used Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | ActiveMQ JVM Memory Usage | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Object Pending FinalizationCount | ActiveMQ JVM Memory Usage | False |
| UTILIZATION\|Process CpuLoad | ActiveMQ OS | False |
| UTILIZATION\|System Cpu Load | ActiveMQ OS | False |
| UTILIZATION\|Consumer Count | ActiveMQ Topic | True |
| UTILIZATION\|Dequeue Count | ActiveMQ Topic | True |
| UTILIZATION\|Enqueue Count | ActiveMQ Topic | True |
| UTILIZATION\|Queue Size | ActiveMQ Topic | True |
| UTILIZATION\|Producer Count | ActiveMQ Topic | False |

## Apache HTTPD Metrics

Metrics are collected for the Apache HTTPD application service.

**NOTE**
Metrics are collected for the Events MPM. Metrics are not collected for the other MPMs.

**Table 77: Apache HTTPD Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Busy Workers | Apache HTTPD | True |
| UTILIZATION\|Bytes Per Req | Apache HTTPD | False |
| UTILIZATION\|Bytes Per Sec | Apache HTTPD | False |
| UTILIZATION\|CPU Load | Apache HTTPD | True |
| UTILIZATION\|CPU User | Apache HTTPD | False |
| UTILIZATION\|Idle Workers | Apache HTTPD | True |
| UTILIZATION\|Request Per Sec | Apache HTTPD | True |
| UTILIZATION\|SCBoard Closing | Apache HTTPD | False |
| UTILIZATION\|SCBoard DNS Lookup | Apache HTTPD | False |
| UTILIZATION\|SCBoard Finishing | Apache HTTPD | False |
| UTILIZATION\|SCBoard Idle Cleanup | Apache HTTPD | False |
| UTILIZATION\|SCBoard Keep Alive | Apache HTTPD | False |
| UTILIZATION\|SCBoard Logging | Apache HTTPD | False |
| UTILIZATION\|SCBoard Open | Apache HTTPD | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|SCBoard Reading | Apache HTTPD | False |
| UTILIZATION\|SCBoard Sending | Apache HTTPD | False |
| UTILIZATION\|SCBoard Starting | Apache HTTPD | False |
| UTILIZATION\|SCBoard Waiting | Apache HTTPD | False |
| UTILIZATION\|Total Accesses | Apache HTTPD | False |
| UTILIZATION\|Total Bytes | Apache HTTPD | True |
| UTILIZATION\|Total Connections | Apache HTTPD | False |
| UTILIZATION\|Uptime | Apache HTTPD | True |
| UTILIZATION\|Asynchronous Closing Connections | Apache HTTPD | False |
| UTILIZATION\|Asynchronous Keep Alive Connections | Apache HTTPD | False |
| UTILIZATION\|Asynchronous Writing Connections | Apache HTTPD | False |
| UTILIZATION\|ServerUptimeSeconds | Apache HTTPD | False |
| UTILIZATION\|Load1 | Apache HTTPD | False |
| UTILIZATION\|Load5 | Apache HTTPD | False |
| UTILIZATION\| ParentServerConfigGeneration | Apache HTTPD | False |
| UTILIZATION\| ParentServerMPMGeneration | Apache HTTPD | False |
| Application Availability | Apache HTTPD | False |

## Oracle DB Metrics

Metrics are collected for the Oracle DB application service.

Oracle DB cannot be activated on Linux platforms.

**Table 78: Oracle DB Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Utilization\|Active Sessions | OracleDB | True |
| Utilization\|Buffer CacheHit Ratio | OracleDB | False |
| Utilization\|Cursor CacheHit Ratio | OracleDB | False |
| Utilization\|Database Wait Time | OracleDB | False |
| Utilization\|Disk Sort persec | OracleDB | False |
| Utilization\|Enqueue Timeouts Persec | OracleDB | False |
| Utilization\|Global Cache Blocks Corrupted | OracleDB | False |
| Utilization\|Global Cache Blocks Lost | OracleDB | False |
| Utilization\|Library CacheHit Ratio | OracleDB | False |
| Utilization\|Logon persec | OracleDB | True |
| Utilization\|Memory Sorts Ratio | OracleDB | True |
| Utilization\|Rows persort | OracleDB | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Utilization\|Service Response Time | OracleDB | False |
| Utilization\|Session Count | OracleDB | True |
| Utilization\|Session Limit | OracleDB | False |
| Utilization\|Shared Pool Free | OracleDB | False |
| Utilization\|Temp Space Used | OracleDB | False |
| Utilization\|Total Sorts persec | OracleDB | False |
| Utilization\|Physical Read Bytes Persc | OracleDB | False |
| Utilization\|Physical Read IO Requests Persc | OracleDB | False |
| Utilization\|Physical Read Total Bytes Persec | OracleDB | False |
| Utilization\|Physical Reads Persec | OracleDB | True |
| Utilization\|Physical Reads Per Txn | OracleDB | False |
| Utilization\|Physical Write Bytes Persc | OracleDB | False |
| Utilization\|Physical Write IO Requests Persc | OracleDB | False |
| Utilization\|Physical Write Total Bytes Persc | OracleDB | False |
| Utilization\|Physical Writes Persc | OracleDB | True |
| Utilization\|Physical Writes Per Txn | OracleDB | False |
| Utilization\|User Commits Percentage | OracleDB | False |
| Utilization\|User Commits Persc | OracleDB | False |
| Utilization\|User Rollbacks Percentage | OracleDB | False |
| Utilization\|User Rollbacks persec | OracleDB | True |
| Utilization\|User Transaction Persec | OracleDB | False |
| Utilization\|Database Time Persc | OracleDB | False |
| Application Availability | Oracle DB | False |

## Cassandra Metrics

Metrics are collected for the Cassandra application service.

**Table 79: Cassandra Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Cache<InstanceName>\|Capacity | Cassandra | False |
| Cache<InstanceName>\|Entries | Cassandra | True |
| Cache<InstanceName>\|HitRate | Cassandra | True |
| Cache<InstanceName>\|Requests | Cassandra | True |
| Cache<InstanceName>\|Size | Cassandra | False |
| ClientRequest<InstanceName>\|Failures | Cassandra | False |
| ClientRequest<InstanceName>\|Latency | Cassandra | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| ClientRequest<InstanceName>\|Timeouts | Cassandra | False |
| ClientRequest<InstanceName>\|Total Latency | Cassandra | False |
| ClientRequest<InstanceName>\|Unavailables | Cassandra | False |
| CommitLog\|Pending Tasks | Cassandra | False |
| CommitLog\|Total Commit Log Size | Cassandra | False |
| Compaction\|Bytes Compacted | Cassandra | False |
| Compaction\|Completed Tasks | Cassandra | False |
| Compaction\|Pending Tasks | Cassandra | False |
| Compaction\|Total Compactions Completed | Cassandra | False |
| Connected Native Clients | Cassandra | False |
| HeapMemoryUsage\|committed | Cassandra | False |
| HeapMemoryUsage\|init | Cassandra | False |
| HeapMemoryUsage\|max | Cassandra | False |
| HeapMemoryUsage\|used | Cassandra | False |
| NonHeapMemoryUsage\|committed | Cassandra | False |
| NonHeapMemoryUsage\|init | Cassandra | False |
| NonHeapMemoryUsage\|max | Cassandra | False |
| NonHeapMemoryUsage\|used | Cassandra | False |
| ObjectPendingFinalizationCount | Cassandra | False |
| Storage\|Exceptions Count | Cassandra | False |
| Storage\|Load Count | Cassandra | False |
| Table<InstanceName>\|Coordinator Read Latency | Cassandra | False |
| Table<InstanceName>\|Live Diskspace Used | Cassandra | False |
| Table<InstanceName>\|Read Latency | Cassandra | False |
| Table<InstanceName>\|Total Diskspace Used | Cassandra | False |
| Table<InstanceName>\|Total Read Latency | Cassandra | False |
| Table<InstanceName>\|Total Write Latency | Cassandra | False |
| Table<InstanceName>\|Write Latency | Cassandra | False |
| ThreadPools<InstanceName>\|Active Tasks | Cassandra | False |
| ThreadPools<InstanceName>\|Currently Blocked Tasks | Cassandra | False |
| ThreadPools<InstanceName>\|Pending Tasks | Cassandra | False |
| Application Availability | Cassandra | False |

## HyperV Metrics

Metrics are collected for the HyperV application service.

**Table 80: HyperV Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| VM:Hyper-V Virtual Machine Health Summary\|Health Critical | HyperV | False |
| VM<instanceName>\|Physical Memory | HyperV | False |
| VM<instanceName>Hv VP 0\|Total Run Time | HyperV | False |
| VM<instanceName>\|Bytes Received | HyperV | False |
| VM<instanceName>\|Bytes Sent | HyperV | False |
| VM<instanceName>\|Error Count | HyperV | False |
| VM<instanceName>\|Latency | HyperV | False |
| VM<instanceName>\|Queue Length | HyperV | False |
| VM<instanceName>\|Throughput | HyperV | False |
| CPU<instanceName>\|Idle Time | HyperV | True |
| CPU<instanceName>\|Processor Time | HyperV | True |
| CPU<instanceName>\|User Time | HyperV | True |
| Disk<instanceName>\|Avg Disk Queue Length | HyperV | False |
| Disk<instanceName>\|Idle Time | HyperV | False |
| Disk<instanceName>\|Read Time | HyperV | True |
| Disk<instanceName>\|Write Time | HyperV | True |
| Process<instanceName>\|Private Bytes | HyperV | False |
| Process<instanceName>\|Processor Time | HyperV | False |
| Process<instanceName>\|Thread Count | HyperV | False |
| Process<instanceName>\|User Time | HyperV | False |
| System\|Processes | HyperV | False |
| System\|Processor Queue Length | HyperV | False |
| System\|System UpTime | HyperV | False |
| Memory\|Available Bytes | HyperV | False |
| Memory\|Cache Bytes | HyperV | False |
| Memory\|Cache Faults | HyperV | False |
| Memory\|Pages | HyperV | False |
| Network<instanceName>\|Packets Outbound Error | HyperV | False |
| Network<instanceName>\|Packets Received Error | HyperV | False |
| Application Availability | HyperV | False |

## MongoDB Metrics

Metrics are collected for the MongoDB application service.

**Table 81: MongoDB Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Active Reads | MongoDB | True |
| UTILIZATION\|Active Writes | MongoDB | True |
| UTILIZATION\|Connections Available | MongoDB | False |
| UTILIZATION\|Connections Total Created | MongoDB | False |
| UTILIZATION\|Current Connections | MongoDB | True |
| UTILIZATION\|Cursor Timed Out | MongoDB | True |
| UTILIZATION\|Deletes Per Sec | MongoDB | False |
| UTILIZATION\|Document Inserted | MongoDB | False |
| UTILIZATION\|Document Deleted | MongoDB | False |
| UTILIZATION\|Flushes Per Sec | MongoDB | False |
| UTILIZATION\|Inserts Per Sec | MongoDB | False |
| UTILIZATION\|Net Input Bytes | MongoDB | False |
| UTILIZATION\|Open Connections | MongoDB | True |
| UTILIZATION\|Page Faults Per Second | MongoDB | False |
| UTILIZATION\|Net Output Bytes | MongoDB | False |
| UTILIZATION\|Queries Per Sec | MongoDB | False |
| UTILIZATION\|Queued Reads | MongoDB | True |
| UTILIZATION\|Queued Writes | MongoDB | True |
| UTILIZATION\|Total Available | MongoDB | False |
| UTILIZATION\|Total Deletes Per Sec | MongoDB | False |
| UTILIZATION\|Total Passes Per Sec | MongoDB | False |
| UTILIZATION\|Total Refreshing | MongoDB | False |
| UTILIZATION\|Updates Per Sec | MongoDB | False |
| UTILIZATION\|Volume Size MB | MongoDB | False |
| Application Availability | MongoDB | False |
| UTILIZATION\|Collection Stats | MongoDB DataBases | False |
| UTILIZATION\|Data Index Stats | MongoDB DataBases | True |
| UTILIZATION\|Data Indexes | MongoDB DataBases | False |
| UTILIZATION\|Data Size Stats | MongoDB DataBases | True |
| UTILIZATION\|Average Object Size stats | MongoDB DataBases | False |
| UTILIZATION\|Num Extents Stats | MongoDB DataBases | False |

## Riak KV Metrics

Metrics are collected for the Riak KV application service.

**Table 82: Riak KV Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|CPU Average | Riak KV | False |
| UTILIZATION\|Memory Processes | Riak KV | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Memory Total | Riak KV | False |
| UTILIZATION\|Node GETs | Riak KV | True |
| UTILIZATION\|Node GETs Total | Riak KV | False |
| UTILIZATION\|Node PUTs | Riak KV | True |
| UTILIZATION\|Node PUTs Total | Riak KV | False |
| UTILIZATION\|PBC Active | Riak KV | True |
| UTILIZATION\|PBC Connects | Riak KV | True |
| UTILIZATION\|Read Repairs | Riak KV | True |
| UTILIZATION\|vNODE Index Reads | Riak KV | True |
| UTILIZATION\|vNODE Index Writes | Riak KV | True |
| Application Availability | Riak KV | False |

## Network Time Protocol Metrics

Metrics are collected for the Network Time Protocol application service.

**Table 83: Network Time Protocol Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| NTPD \|delay | Network Time Protocol | True |
| NTPD \| jitter | Network Time Protocol | True |
| NTPD \| offset | Network Time Protocol | True |
| NTPD \| poll | Network Time Protocol | False |
| NTPD \| reach | Network Time Protocol | True |
| NTPD \| when | Network Time Protocol | False |
| Application Availability | Network Time Protocol | False |

## WebSphere Metrics

Metrics are collected for the WebSphere application service.

**Table 84: WebSphere Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Thread Pool\|Active Count\| Current | Thread Pool | False |
| Thread Pool\|Active Count\| High | Thread Pool | False |
| Thread Pool\|Active Count\| Low | Thread Pool | False |
| Thread Pool\|Active Count\| Lower | Thread Pool | False |
| Thread Pool\|Active Count\| Upper | Thread Pool | False |
| JDBC\|Close Count | JDBC | False |
| JDBC\|Create Count | JDBC | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JDBC\|JDBC Pool Size\| Average | JDBC | False |
| JDBC\|JDBC Pool Size\| Current | JDBC | False |
| JDBC\|JDBC Pool Size\| Lower | JDBC | False |
| JDBC\|JDBC Pool Size\| Upper | JDBC | False |
| Garbage Collection<InstanceName>\| Total Collection Count | WebSphere | False |
| Garbage Collection<InstanceName>\| Total Collection Time | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Number of Object Pending Finalization Count | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Committed Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Initial Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Maximum Memory | WebSphere | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Pool<InstanceName>\|Peak Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Pool<InstanceName>\|Usage\|Used Memory | WebSphere | False |
| Process Cpu Load | WebSphere | False |
| System Cpu Load | WebSphere | False |
| System Load Average | WebSphere | False |
| Application Availability | WebSphere | False |

## Java Application Metrics

Metrics are collected for the Java application service.

**Table 85: Java Application Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Java Application | False |
| Buffer Pool<InstanceName>\|Memory Used | Java Application | False |
| Buffer Pool<InstanceName>\|Total Capacity | Java Application | False |
| Class Loading\|Loaded Class Count | Java Application | True |
| Class Loading\|Total Loaded Class Count | Java Application | False |
| Class Loading\|Unloaded Class Count | Java Application | False |
| Garbage Collection<InstanceName>\|Total Collection Count | Java Application | False |
| Garbage Collection<InstanceName>\|Total Collection Time | Java Application | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\|Initial Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\|Maximum Memory | Java Application | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Heap Memory Usage\| Used Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Committed Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Initial Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Maximum Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Used Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\| Committed Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Initial Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\| Maximum Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Used Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Java Application | False |
| JVM Memory\|Object Pending Finalization Count | Java Application | False |
| Uptime | Java Application | True |
| Threading\|Thread Count | Java Application | True |
| Process CPU Usage % | Java Application | False |
| System CPU Usage % | Java Application | False |
| System Load Average % | Java Application | False |

**Remote Check Metrics**

Metrics are collected for object types such as HTTP, ICMP, TCP, and UDP.

## HTTP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for HTTP remote checks.

HTTP Metrics

**Table 86: HTTP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Content Length | False |
| Response Code | False |
| Response Time | True |
| Result Code | False |

## ICMP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for the ICMP object type.

**Table 87: ICMP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Average Response Time | True |
| Packet Loss (%) | False |
| Packets Received | False |
| Packets Transmitted | False |
| Result Code | False |

## TCP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for the TCP object type.

**Table 88: TCP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Response Time | True |
| Result Code | False |

## UDP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for the UDP object type.

**Table 89: UDP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Response Time | True |
| Result Code | False |

### Linux Process Metrics

Metrics are collected for Linux services.

**Table 90: Linux Process Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| AVAILABILITY\|Resource Availability | Processes | False |
| UTILIZATION\|Memory Usage (%) | Processes | False |
| UTILIZATION\|CPU Usage (%) | Processes | False |
| UTILIZATION\|Number of Processes | Processes | False |

### Windows Service Metrics

Metrics are collected for Windows services.

**Table 91: Windows Service Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| AVAILABILITY\|Resource Availability | Services | False |
| UTILIZATION\|Memory Usage(%) | Services | False |
| UTILIZATION\|CPU Usage(%) | Services | False |

### OS and Application Monitoring Properties

Properties are collected for operating systems, application services, remote checks, Linux processes, and Windows services which can be used to create reports, views, and dashboards.

### Guest Information Properties

VMware Aria OperationsVMware Cloud Foundation Operations displays the following guest information properties for all objects created by the OS and Application Monitoring management pack.

- Guest Info
  - Hostname
  - IP
  - OS Name
  - OS Version
  - Telegraf Version

Other properties of operating systems and application services are available under **Properties › Tags**.

### Troubleshooting

## Troubleshooting Agent Installation

### Agent Install Failure Because of the vCenter User Permissions

Guest operation privileges are required to install agents on end point virtual machines.

Agent installation fails with the following error message if there are no guest operation privileges:

```
vCenter adapter user is missing either of the following guest operations privileges –
execute, modify, query
```

1. Verify that you have configured a vCenter adapter.
2. The vCenter user account with which the vCenter adapter is configured in VMware Aria OperationsVMware Cloud Foundation Operations, should have the following permissions: **Guest operation modifications**, **Guest operation program execution**, and **Guest operation queries**.

### Agent Install Failure Because NTP is Not in Sync

After you install or upgrade to the latest version of cloud proxy, you must set up accurate timekeeping as part of the deployment. If the time settings between cloud proxy and VMware Aria OperationsVMware Cloud Foundation Operations are not synchronized, you face agent installation and metric collection issues. Ensure time synchronization between the endpoint VMs, vCenter Server, ESX Hosts, cloud proxy and VMware Aria OperationsVMware Cloud Foundation Operations using the Network Time Protocol (NTP).

- Agent installation fails
- Log in to cloud proxy and run the following command to stop the NTP daemon:

  ```
  systemctl stop ntpd
  ```

- Run the following command to update the time immediately from an NTP server:

  ```
  ntpdate time.vmware.com
  ```

  **NOTE**
  Replace `time.vmware.com` with a suitable time server setting. You can use the FQDN or IP of the time server.

- Enter the following command to start the NTP daemon:

  ```
  systemctl start ntpd
  ```

  **NOTE**
  The system time takes about five minutes to sync with the NTP server time.

### Agent Install Fails on a Linux End Point

Install of an agent on a Linux end point fails for a non-root user with a specific set of privileges.

Agent installation fails with the following error if the **`tty`** command is not added:

```
Bootstrap Failed for VM <VM ID> with error message:{ "status":"FAILED", "data":
[ { "status":"FAILED", "message":"Failed - install - passwordless sudo access is required
for the user <Install Username> on the command mkdir. [sudo: sorry, you must have a tty to
run sudo]", "stage":"0" } ], "currentstage":"0", "totalstages":"0" }
```

The following are minimal necessary permissions of the user to install agents and should be mentioned in the `sudoers` file:

- The following are minimal necessary permissions of the user to install agents and should be mentioned in the `sudoers` file. For example, for a user called **telegrafinstall**, you can find the sudoers file in the **/etc/sudoers** file or in the folder **/etc/sudoers.d/:**. Add these lines to `/etc/sudoers`, if you have not added them.

```
Defaults:telegrafinstall !requiretty



Cmnd_Alias ARC_INSTALL_USER_COMMANDS=/usr/bin/cp*,/bin/cp*,/usr/bin/mkdir*,/bin/
mkdir*,/usr/bin/chmod*,/bin/chmod*,/opt/vmware/ucp/bootstrap/uaf-bootstrap.sh,/opt/
vmware/ucp/ucp-minion/bin/ucp-minion.sh



telegrafinstall ALL=(ALL)NOPASSWD: ARC_INSTALL_USER_COMMANDS
```

### Agent Install Fails on Windows When UAC is Deactivated

Install of the agent fails even when UAC is deactivated.

1. Make sure UAC (previously known as LUA) was deactivated on Windows in the following way:
   a) In the registry path `HKLM:\Software\Microsoft\Windows\CurrentVersion\Policies\System`, set the value for the key `EnableLUA` to **0**.
   b) You must reboot the machine for the changes to take effect.

### Agent Install Fails on Windows with a Permission Denied Error

In Windows, during bootstrap, when the Telegraf folder is renamed to `ucp-telegraf`, it can result in a failure because of a permission error.

Sometimes, there are certain antiviruses running, which prevent the application from renaming or modifying the directory or files. In such a situation, the following error message is displayed:

**`Install telegraf [unable to install telegraf due to system error : [WinError 5] Access is denied: 'C:\\VMware\\UCP\\ucp-telegraf']"}].`**

1. Deactivate the antivirus and then proceed with bootstrapping.

### Agent Install Does Not Progress

During agent install, new tasks do not progress beyond the Starting phase, from recent tasks. Adapter logs are not written. Verify that the adapter instance in cloud proxy is in a Data Collecting state. If not, restart the adapter instance from the VMware Cloud Foundation Operations user interface. From the left menu, navigate to **Operations › Configurations › Adapter Instances › VMware Aria Operations Application Management Adapter Instance**. From the **Objects** tab in the right pane, select the adapter instance and click **Stop Collecting** and then **Start Collecting**.

### Agent Install Fails Without an Error Message

Agent install fails without any error message, in the user interface

The `uaf_bootstrap.log` at the endpoint VM displays the following log message: `'findstr' is not recognized as an internal or external command`.
This happens because `C:\Windows\System32` is not available in the environment variable *`PATH`*.
Add `C:\Windows\System32` to the environment variable *`PATH`*.

### Script Download Fails on a Windows Platform

When a file is downloaded from cloud proxy to a Windows end point, it could fail due to security protocols.

Script download fails on a Windows platform with the following message:

```
The request was aborted: Could not create SSL/TLS secure channel.
```

There are three kinds of PowerShell scripts hosted in cloud proxy that can be downloaded and executed at the Windows end point VMs for different purposes:

- To install custom Telegraf using a script (download.ps1).
- To install custom Telegraf on a physical server (unmanagedagent_setup_sample.ps1).
- To configure open source Telegraf on managed or unmanaged VMs (open_source_telegraf_monitor.ps1).

Ignore the **`ServerCertificateValidationCallback`** using the following command.

```
if (-not
([System.Management.Automation.PSTypeName]'ServerCertificateValidationCallback').Type)

{

$certCallback = @"

    using System;

    using System.Net;

    using System.Net.Security;

    using System.Security.Cryptography.X509Certificates;

    public class ServerCertificateValidationCallback

    {

        public static void Ignore()

        {

            if(ServicePointManager.ServerCertificateValidationCallback ==null)

            {

                ServicePointManager.ServerCertificateValidationCallback +=

                    delegate

                    (

                        Object obj,
```

```
                        X509Certificate certificate,
                        X509Chain chain,
                        SslPolicyErrors errors
                    )
                    {
                        return true;
                    };
            }
        }
    }
"@
    Add-Type $certCallback
 }
[ServerCertificateValidationCallback]::Ignore()
[Net.ServicePointManager]::SecurityProtocol = [Net.SecurityProtocolType]::Tls12
```

After downloading and executing the required script, **`ServerCertificateValidationCallback`** can be activated.

### Troubleshooting Plugin Related Failures

### Unable to Activate a Plugin

Unable to activate a plugin with the same fields until the plugin configuration is deleted.

An error message is displayed in the user interface of VMware Aria OperationsVMware Cloud Foundation Operations that states the following:

```
Failed to update resource: Resource with same key already exists
```

1. Manually delete the existing plugin configuration and then continue with the activation of the plugin. If the problem persists, delete the corresponding resource from the inventory.

### RabbitMQ Plugin Error

You might receive an error while monitoring the RabbitMQ plugin.

While monitoring the RabbitMQ plugin, an error might occur if you upgrade to VMware Aria OperationsVMware Cloud Foundation Operations, which has a Telegraf version of 1.19. The following message is displayed:

`getting "/api/federation-links" failed: 404 Not Found`
The Federation plugins in the RabbitMQ setup are not activated.

1. Activate the Federation plugins in the RabbitMQ setup by running the following commands in the RabbitMQ VM:

   ```
   rabbitmq-plugins enable rabbitmq_federation
   ```

   ```
   rabbitmq-plugins enable rabbitmq_federation_management
   ```

### Troubleshooting Metric Collection

### Troubleshoot Agent Installation and Metric Collection Issues

If the time settings between cloud proxy and VMware Aria OperationsVMware Cloud Foundation Operations are not synchronized, you might face agent installation and metric collection issues. Eventually, you might not see any metrics in the VMware Aria OperationsVMware Cloud Foundation Operations dashboards.

You might notice the following issues in VMware Aria OperationsVMware Cloud Foundation Operations:
- You cannot install an agent in the Windows and Linux target VMs.

Time synchronization is a prerequisite of the TLS/SSO communication between client and server.

If the VMware Aria OperationsVMware Cloud Foundation Operations and cloud proxy are not time synchronized, the test connection fails while configuring cloud proxy in VMware Aria OperationsVMware Cloud Foundation Operations.

If the Windows and Linux target VMs are not time synchronized with VMware Aria OperationsVMware Cloud Foundation Operations, communication between cloud proxy and agents will break after installing the agents. Hence monitored metrics are not sent to VMware Aria OperationsVMware Cloud Foundation Operations. Alternatively, stop and restart the agent to resolve this issue.

1. Check the VMware Aria OperationsVMware Cloud Foundation Operations support bundle in the following path: `COLLECTOR/adapters/APPOSUCPAdapter/` for errors.
2. Check the cloud proxy support bundle, *`ucpapi.log`*, for errors.
3. Ensure time synchronization between cloud proxy, VMware Aria OperationsVMware Cloud Foundation Operations and the Windows and Linux target VMs.
4. To start and restart the agent, see Additional Operations from the Manage Telegraf Agents Page .

### Troubleshooting Content Upgrade for an End Point

Content upgrade for an end point fails with the following error:

**`Timeout Error. Please retry the action after some time.`**
Sometimes content upgrade for an end point fails because of a timeout in the cloud proxy.

1. Retrigger content upgrade for the end point to resolve the issue.

### Telegraf Agent Related Actions Fail after Cloud Proxy is Restarted or Upgraded

After an upgrade of vRealize Operations/vRealize Operations Cloud and cloud proxy from 8.4 to a later releases or if you restart cloud proxy, Telegraf agent management related actions could fail with the error message `Connect to Salt Master`.
The saltmaster docker container inside cloud proxy does not run properly causing all Telegraf related actions to fail.

1. SSH to the cloud proxy VM.
2. Verify if the salt process count by command is: **`ps -ef | grep salt`**.
3. If the salt process count is less than 15, run the following command: **`/rpm-content/ucp/subsequentboot.sh`**.

   You can view the log from the following location: `/opt/vmware/var/log/ucp-subsequentboot`.

### Troubleshooting Using Support Bundles

Support bundles are required to troubleshoot problems related to application monitoring. For Linux and Windows end point VMs, run the specified command and access the support bundle.

#### For End Point VMs

1. Log in to the end point.
2. Run the following commands based on the end point VM's operating system type:
   **For Linux End Point VMs**
   ```
   /opt/vmware/ucp/ucp-minion/bin/ucp-minion.sh --config /opt/vmware/ucp/salt-minion/
   etc/salt/grains --action gen_support_bundle --log_level INFO
   ```

   The support bundle is generated and placed as a ZIP file in the `/opt/vmware/ucp/support-bundle-endpoints/` directory.

   **For Windows End Point VMs**
   ```
   C:\VMware\UCP\ucp-minion\bin\ucp-minion.bat --config C:\VMware\UCP\salt\conf\grains
   --action gen_support_bundle --log_level INFO
   ```

   The support bundle is generated and placed as a ZIP file in the `%SystemDrive%\VMware\UCP\support-bundle-endpoints\` directory.

### Objects Do Not Receive Data from the Adapter Instance

In case of a HA failover, you see the following alert: `Objects are not receiving data from the adapter instance`. This is an adapter-specific alert that does not determine whether the object is receiving data from any of the adapters. It is recommended that you deactivate this alert.
1. Create a new policy that inherits the AppOS adapter instance's pre-applied policy. From the left menu, click **Operations › Configurations**, and then from the left panel, click the **Policy Definition** tile.
2. Click **Add** to add a policy, or select the policy and click **Edit Policy** to edit an existing policy.

   You can add and edit policies and remove certain policies. You can use the Base Settings policy or the Default Policy as the root policy for the settings in other policies that you create. You can set any policy to be the default policy.

3. In the **Create Policies** workspace, assign a name to the policy and enter the description.
   a) Give the policy a meaningful name and description so that all users know the purpose of the policy.
   b) From the **Inherit From** drop-down option, select the policy to use as a baseline to define the settings for your new local policy.

4. Click **Create Policy**. The **Create Policies** workspace provides options to customise your policy.
5. Click the **Alerts and Symptoms** tile and search for the alert called `Objects are not receiving data from adapter instances` from the policy created in step 1.
6. Click on the alert and deactivate it.
7. Apply the policy created in step 1 to the AppOS adapter instances.

## Monitoring Application Services and Operating Systems using Open Source Telegraf

You can use open source Telegraf that runs on a vCenter VM, an EC2 AWS instance, or an Azure VM to send metrics to VMware Aria OperationsVMware Cloud Foundation Operations using a helper script or by providing specific configurations required on the end point to post metrics to cloud proxy. The helper script adds necessary configurations in the Telegraf configuration that is directly identified by VMware Aria OperationsVMware Cloud Foundation Operations.

You can monitor application services and operating systems that are supported by VMware Aria OperationsVMware Cloud Foundation Operations and you can also monitor unsupported application services. For a list of supported application services, see Supported Application Services .

Open source Telegraf is supported on Linux and Windows platforms.

**Notice:** VMware supports only the installation of open source Telegraf because there is a configuration required to point it to the VMware Aria OperationsVMware Cloud Foundation Operations cloud proxy. After the initial configuration is complete, and we have made sure that the metrics are received consistently for a couple of end points, and there are no common patterns of failures observed on VMware Aria OperationsVMware Cloud Foundation Operations, no further support is offered from VMware. Any other queries on the agent has to go through the Telegraf open source community. Issues related to Telegraf can be raised through: https://github.com/influxdata/telegraf/issues.

**Points to Note**

As you can monitor both supported and unsupported application services using open source Telegraf, keep in mind the following points.

- There are no alerts generated for non-supported application services.
- You cannot add metrics for supported application services.
- Application service objects that are not supported are named as follows: *<application service name_Generic>*
- Metrics for unsupported application services are displayed without proper categorization and appears as is from Telegraf.
- Metrics for unsupported application services are not localized and are provided only in English.
- If there are multiple instances of unsupported application services, to distinguish between them in the user interface, add the tag 'identifier' and provide a unique tag value in the `telegraf.conf` file. Example, [inputs.mongodb.tags] identifier="1".
- If multiple instances of the same supported application service is monitored by open source Telegraf, in the user interface, the display name is the same for all instances of the application service. To differentiate between the instances of the same application service, edit the display name from the user interface and provide a specific name.
- Inactivated application services with a managed agent are displayed as activated in the **Manage Telegraf Agents** page. Uninstall managed agents before you use open source Telegraf.
- Telegraf version 1.19.0 is not supported when you use open source Telegraf.

### Monitoring Application Using Open Source Telegraf

You can use open source Telegraf to monitor application services and operating systems on a or Linux platform or on a Windows platform.

### Install and Configure Open Source Telegraf

Install and configure open source Telegraf to monitor your applications.

Telegraf uses input plugins (where the metrics come from) and output plugins (where the metrics go) in the configuration files. You can see all the supported plugins at Plugin Directory. Input and output plugins should be written in the Telegraf configuration file and configuration directory. The Telegraf configuration file (telegraf.conf) lists all available Telegraf plugins.
Telegraf uses the `--config` flag to specify the configuration file location and `--config-directory` flag to include files ending with `.conf` in the specified directory. On most Linux systems, the default locations are `/etc/telegraf/telegraf.conf` for the main configuration file and `/etc/telegraf/telegraf.d` for the directory of configuration files. For Windows platforms, it will be in the location where the `telegraf zip` is extracted.
For more information, see the Telegraf documentation, Get Started and Configuration Options. After you have downloaded and installed Telegraf, you are ready to begin collecting and sending data. To collect and send data Telegraf should be configured. Follow the steps below:

1. Install open source Telegraf on the end point. If you have an instance installed, you can skip this step. To download and install a new instance of Telegraf, see the official documentation and search for the corresponding OS version from https://www.influxdata.com/time-series-platform/telegraf/ and https://portal.influxdata.com/downloads/.

   On a Windows platform only, after downloading and extracting Telegraf files, besides `telegraf.exe` and `telegraf.conf` files, create a folder with the name `telegraf.d` that you use to run the helper script and monitor applications.

2. Run the helper script to configure Telegraf to send data to cloud proxy.

   For more information, see Monitoring Applications Using Open Source Telegraf on a Linux Platform and Monitoring Applications Using Open Source Telegraf on a Windows Platform.

   After you run the helper script, a `cloudproxy-http.conf` file is created and the output plugin is added to the file with the properties necessary for sending data and the input plugins are updated for OS metrics. For OS metrics configuration template, see Telegraf Configuration Details for Operating Systems. For information about the `cloudproxy-http.conf` configuration template, see Sample Configurations.

3. Using open source Telegraf, you can collect metrics from different types of application services.
   - If an application service (curated plugin) that is supported by Application Monitoring in VMware Aria OperationsVMware Cloud Foundation Operations is running on the end point and you want to monitor it, update the Telegraf configuration file or directory with necessary inputs for Telegraf.
     For a list of supported application services (curated plugins), see Supported Applications. For the list of configurations, see Telegraf Configuration Details for Supported Application Services.
   - For unsupported (non-curated plugins) application services, update the Telegraf configuration file or directory with the necessary inputs for Telegraf. For more information, see the Telegraf documentation Plugin Directory. For example, if you want to collect data for Ethernet device stats (device name is eth0) you can create the `ethools.conf` file in the Telegraf configuration directory and add the following content (for more information, see the documentation about Ethtool Input Plugin):

     ```
     [[inputs.ethtool]]

     # List of interfaces to pull metrics for

      interface_include = ["eth0"]
     ```

   - Whenever the `telegraf-utils` script is executed with Open Source Telegraf, the `cloudproxy-http.conf` file will be overwritten in the `telegraf.d` folder. You must maintain a different config file (for example, `postgres.conf`) for other custom input plugins to retain the configuration.

4. Start or restart Telegraf.

### Monitoring Applications using Open Source Telegraf on a Linux Platform

Use the helper script to monitor applications and operating systems on a Linux platform using open source Telegraf.

- Install the jq package. For more information, see the official documentation for jq from https://stedolan.github.io/jq/download/.
- Ensure that VMTools version >=10.2 if a vCenter VM is monitored by VMware Aria OperationsVMware Cloud Foundation Operations. **`vmware-toolbox-cmd`** is necessary only for vCenter VMs.
- Ensure that you have permissions to run the following commands and use the following packages:
  - Commands: /bin/bash, awk, sed, vmware-toolbox-cmd, jq, curl
  - Packages: coreutils (chmod, chown, cat), net-tools (ip, /dev/tcp, curl, wget)
- Only IPv4 is supported at present for cloud proxy.
- Ensure that Internet is activated.
- Ensure that cloud proxy is up and online in VMware Aria OperationsVMware Cloud Foundation Operations. In the case of collector groups, ensure that at least one cloud proxy in the application monitoring high availability activated collector group is up and online in VMware Aria OperationsVMware Cloud Foundation Operations.
- Verify that unzip is at 6.0-20.el7 or above.
- The uuidgen package must exist on the endpoint VM/physical server.
- The uuidgen package must exist on the vCenter VMs, AWS EC2 instances, Azure VMs, or the physical server.
- The endpoint VM/physical server must have access to port 8443 and 443 of cloud proxy or the virtual IP of the application monitoring high availability activated collector group.
- vCenter VMs, AWS EC2 instances, Azure VMs, or the physical server must have access to port 8443 and 443 of cloud proxy or the virtual IP of the application monitoring high availability activated collector group.

The helper script is tested only on the following operating systems:

- CentOS 7.x and CentOS 8.x
- RHEL 7.x and RHEL 8.x
- SUSE 12.x and SUSE 15.x
- OEL7.x and OEL 8.x
- Ubuntu 16.x, Ubuntu 18.x, Ubuntu 20.x, and Ubuntu 22.x
- VMware Photon Linux

1. Download the helper script from cloud proxy located at `https://<CloudProxy-IP>/downloads/salt/telegraf-utils.sh`.

   **NOTE**
   Use the relevant cloud proxy **IP address** for <CloudProxy-IP> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

2. Navigate to the directory where the script is downloaded.
3. Activate execution permission of the script for Linux VM.

   ```
   chmod +x telegraf-utils.sh
   ```

4. Run the helper script to update Telegraf configurations.

   ```
   telegraf-utils.sh  opensource -c cloud_proxy_ip_or_collector_group_name -t token -d
   telegraf_conf_dir -e telegraf_bin_path -v vmwarearíaoperations_ip_or_fqdn[-g
   gateway_url -a csp_auth_url]


   Description of arguments:
   ```

```
      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or
Collector Group Name

       -t : [REQUIRED] token - CSP Refresh Token of the user/account. For getting a
new token,

          follow - User/Organization Settings >> My Account >> API Tokens >> Generate
a New API Token.

          Example: gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq
 -t : [REQUIRED] token - Auth Token of the user/account. For getting a new token:

          (https://<VMwareAriaOperations_IP>/suite-api/ or curl -ks -X POST https://
<VMwareAriaOperations_IP>/suite-api/api/auth/token/acquire -H \"Content-Type:
application/json\" -H \"Accept: application/json\" -d \"{\"username\":
\"<VMwareAriaOperations_USER>\",\"password\":
\"<VMwareAriaOperations_USER_PASSWORD>\"}\")

      -d : [REQUIRED] telegraf_conf_dir - Telegraf configuration directory and it is
required argument. ex: /etc/telegraf/telegraf.d

      -e : [REQUIRED] telegraf_bin_path - Path of telegraf executable ex: /usr/bin/
telegraf

      -v : [REQUIRED] vmwareariaoperations_ip_or_fqdn - IP/FQDN of VMware Aria
Operations master node and required for on-prem.-g : [OPTIONAL] [SAAS-SPECIFIC]
gateway_url - argument to override default VMware Aria Operations SaaS gateway URL

      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URLExample: /bin/bash telegraf-utils.sh opensource -t
gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq -v 10.10.10.100 -c
10.10.10.101 -d /etc/telegraf/telegraf.d -e /usr/bin/telegrafExample: /bin/bash
telegraf-utils.sh opensource -t
gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq -c 10.10.10.101 -d /
etc/telegraf/telegraf.d -e /usr/bin/telegraf
```

**NOTE**
After you run the helper script, ensure that the respective configurations are set correctly in the given config directory (`-d` option) path with the name `cloudproxy-http.conf`. See Sample Configurations for more details. For managed vCenter VMs, AWS EC2 instances, or Azure VMs, you might see unmanaged configurations, because of one of the following reasonsFor managed vCenter VMs you might see unmanaged configurations, because of one of the following reasons:

- vCenter VM details are not available in VMware Aria OperationsVMware Cloud Foundation Operations by the vCenter adapter. Wait for a minimum of one to two collection cycles after configuring the VMware Aria OperationsVMware Cloud Foundation OperationsvCenter cloud accounts.
- vCenter VMs, AWS EC2 instances, or Azure VMs details are not available in VMware Aria OperationsVMware Cloud Foundation Operations by the vCenter, AWS, or Azure adapters correspondingly. Wait for a minimum of one to two collection cycles after configuring the VMware Aria OperationsVMware Cloud Foundation OperationsvCenter cloud accounts.
- An incorrect `SAAS_REFRESH_TOKEN`.
- An incorrect `AUTHENTIFICATION_TOKEN or vROps_IP`.

**NOTE**
By default, the InfluxDB output plugin is active in the `telegraf.conf` file and data is sent to the influxdb server so that you do not get multiple warning messages in the logs about the lack of configured influxdb server comment, the "[[outputs.influxdb]]" line should be commented. The following warning message is displayed:**`W! [outputs.influxdb] When writing to [http://localhost:8086]: database "telegraf" creation failed: Post "http://localhost:8086/query": dial tcp [::1]:8086: connect: connection refused`**
Example: `#[[outputs.influxdb]]`

**NOTE**
Ensure that the input plugins in the `telegraf.conf` file are related to the corresponding operating system. See Telegraf Configuration Details for Operating Systems.

**NOTE**
The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

5. Restart the Telegraf service.

```
systemctl restart telegraf

or

/usr/bin/telegraf -config /etc/telegraf/telegraf.conf -config-directory /etc/telegraf/telegraf.d
```

**Managed VM object hierarchy**:
- If a vCenter VM is monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects fall under the respective **VM › OS object › 'application service' instance**.
- If an Azure VM is monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects fall under the respective **Azure VM › OS object › 'application service' instance**.
- If an AWS EC2 instance is monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects fall under the respective **AWS EC2 instance › OS object › 'application service' instance**.

**Machines not monitored by VMware Aria Operations**: If a vCenter VM, an AWS EC2 instance, an Azure VM, or a physical server is not monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects fall under **Operations › Configurations › Inventory Management › Endpoint › OS Object › 'application service' instance**.If a vCenter VM or a physical server is not monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects fall under **Operations › Configurations › Inventory Management › Endpoint › OS Object › 'application service' instance**.

**View Data Collected**: Data is collected and appears in the **Manage Telegraf Agents** page. To view the details, you can filter by **Agent Type › Open source agent** from the **Manage Telegraf Agents** page.

### Monitoring Applications using Open Source Telegraf on a Windows Platform

Use the helper script to monitor applications and operating systems on a Windows platform using open source Telegraf.

- Verify that Windows PowerShell is at 4.0 or above.
- Ensure that VMTools version >=10.2 if a vCenter VM is monitored by VMware Aria OperationsVMware Cloud Foundation Operations. **`vmware-toolbox-cmd`** is necessary only for vCenter VMs.

- Only IPv4 is supported at present for cloud proxy.
- Ensure that Internet is enabled.
- Ensure that cloud proxy is up and online in VMware Aria OperationsVMware Cloud Foundation Operations. In the case of collector groups, ensure that at least one cloud proxy in the application monitoring high availability activated collector group is up and online in VMware Aria OperationsVMware Cloud Foundation Operations.
- The endpoint VM/physical server must have access to port 8443 and 443 of cloud proxy or the virtual IP of the application monitoring high availability activated collector group.
- vCenter VMs, AWS EC2 instances, Azure VMs, or the physical server must have access to port 8443 and 443 of cloud proxy or the virtual IP of the application monitoring high availability activated collector group.

The helper script is tested only on Windows Server 2012, 2012 R2, 2016, 2019, and 2022.

1. Download the helper script from cloud proxy located at `https://<CloudProxy-IP>/downloads/salt/telegraf-utils.ps1`.

   **NOTE**
   Use the relevant cloud proxy **IP address** for <CloudProxy-IP> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP.

   If the script download fails with the following message: `The request was aborted: Could not create SSL/TLS secure channel`, follow the steps mentioned in Script Download Fails on a Windows Platform.

2. Navigate to the directory where the script is downloaded.
3. Run the helper script to update Telegraf configurations.

```
telegraf-utils.ps1  opensource -c cloud_proxy_ip_or_collector_group_name -t token -d
telegraf_conf_dir -e telegraf_bin_path -v vmwareariaoperations_ip_or_fqdn[-g
gateway_url -a csp_auth_url]



Description of arguments:

      -c : [REQUIRED] cloud_proxy_ip_or_collector_group_name - Cloud Proxy IP or
Collector Group Name

      -t : [REQUIRED] token - [REQUIRED] - CSP Refresh Token of the user/account. For
getting a new token, follow -

           User/Organization Settings >> My Account >> API Tokens >> Generate a New
API Token.

           Example:
gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq-t : [REQUIRED] token
- Auth Token of the user/account. For getting a new token:

           (https://<VMwareAriaOperations_IP>/suite-api/ or curl -ks -X POST https://
<VMwareAriaOperations_IP>/suite-api/api/auth/token/acquire -H \"Content-Type:
application/json\" -H \"Accept: application/json\" -d \"{\"username\":
\"<VMwareAriaOperations_USER>\",\"password\":
\"<VMwareAriaOperations_USER_PASSWORD>\"}\")

      -d : [REQUIRED] telegraf_conf_dir - Telegraf configuration directory and it is
required argument. ex: C:\Telegraf\telegraf-1.20.4\telegraf.d

      -e : [REQUIRED] telegraf_bin_path - Path of telegraf executable ex: C:
\Telegraf\telegraf-1.20.4\telegraf.exe

      -v : [REQUIRED] vmwareariaoperations_ip_or_fqdn - IP/FQDN of VMware Aria
Operations master node and required for on-prem.-g : [OPTIONAL] [SAAS-SPECIFIC]
gateway_url - argument to override default VMware Aria Operations SaaS gateway URL
```

```
      -a : [OPTIONAL] [SAAS-SPECIFIC] csp_auth_url - argument to override default CSP
authentication URLExample: .\telegraf-utils.ps1 opensource -c 10.10.10.101 -t
41ef6601-6da4-4757-a51d-cbc08dd77355::4398b23d-e388-496e-ae91-bc04d5735345 -v
10.10.10.100 -d C:\Telegraf\telegraf-1.20.4\telegraf.d -e C:
\Telegraf\telegraf-1.20.4\telegraf.exeExample: .\telegraf-utils.ps1 opensource -c
10.10.10.101 -t gi7lwabjnvdfiawt4watzksuol8sywrjvg8kabh31mx9x1guepgyhycyx61dqrpq -d
C:\Telegraf\telegraf-1.20.4\telegraf.d -e C:\Telegraf\telegraf-1.20.4\telegraf.exe
```

**NOTE**
After you run the helper script, ensure that the respective configurations are set correctly in the given config directory (`-d` option) path with the name `cloudproxy-http.conf`. See Sample Configurations for more details. For managed vCenter VMs, AWS EC2 instances, or Azure VMs, you might see unmanaged configurations, because of one of the following reasons:For managed vCenter VMs you might see unmanaged configurations, because of one of the following reasons:

- vCenter VMs are not available in VMware Aria OperationsVMware Cloud Foundation Operations by the vCenter. Wait for a minimum of one to two collection cycles after configuring the VMware Cloud Foundation OperationsvCenter cloud accounts.
- vCenter VMs, AWS EC2 instances, or Azure VMs details are not available in VMware Aria OperationsVMware Cloud Foundation Operations by the vCenter, AWS, or Azure adapters correspondingly. Wait for a minimum of one to two collection cycles after configuring the VMware Cloud Foundation OperationsvCenter cloud accounts.
- An incorrect `SAAS_REFRESH_TOKEN`.
- An incorrect `AUTHENTIFICATION_TOKEN or vROps_IP`.

**NOTE**
- Do not use a space in the configuration path. Paths with spaces can be passed as a short name notation, such as `c:\PROGRA~1` for `c:\Program Files`.

**NOTE**
Ensure that the input plugins in the `telegraf.conf` file are related to the corresponding operating system. See Telegraf Configuration Details for Operating Systems.

**NOTE**
The default gateway URL is `https://www.mgmt.cloud.vmware.com/vrops-cloud` and the default authentication URL is `https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/authorize`. If the authentication URL and gateway URL are different from the default settings, provide the appropriate arguments (-g and -a).

4. Restart the Telegraf service.

```
telegraf.exe --config telegraf.conf --config-directory telegraf.d
```

Or you can make Telegraf a Windows service.

```
<Telegraf_executable_path> --config <Telegraf_config_file_path> --config-directory
<Telegraf_config_directory_path> --service install net start telegraf
```

For example:

```
& 'C:\Telegraf\telegraf-1.20.4\telegraf.exe'  --config 'C:
\Telegraf\telegraf-1.20.4\telegraf.conf' --config-directory 'C:
\Telegraf\telegraf-1.20.4\telegraf.d'  --service install

net start telegraf
```

**Managed VM object hierarchy**:
- If a vCenter VM is monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects will fall under the respective **VM › OS object › 'application service' instance**
- If an Azure VM is monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects will fall under the respective **Azure VM › OS object › 'application service' instance**.
- If an AWS EC2 instance of the VM is monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects will fall under the respective **AWS EC2 instance › OS object › 'application service' instance**

**Machines not monitored by VMware Aria Operations**: If a vCenter VM, an AWS EC2 instance, an Azure VM, or a physical server is not monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects will fall under **Operations › Configurations › Inventory Management › Endpoint › OS Object › 'application service' instance**.If a vCenter VM or a physical server is not monitored by VMware Aria OperationsVMware Cloud Foundation Operations, then the operating system and application objects will fall under **Operations › Configurations › Inventory Management › Endpoint › OS Object › 'application service' instance**.

**View Data Collected**: Data is collected and appears in the **Manage Telegraf Agents** page. To view the details, you can filter by **Agent Type › Open source agent** from the **Manage Telegraf Agents** page.

### Telegraf Configuration Details for Operating Systems

#### Linux Operating Systems

To collect Linux OS related metrics and properties as it is in the managed Telegraf agent installation flow and to have localized object types, names, alerts, symptoms, metrics and properties, and so on, the `telegraf.conf` file (in case of installed telegraf it is located in `/etc/telegraf/telegraf.conf`) should be customized and have the following content:

```
# Read metrics about cpu usage

[[inputs.cpu]]

  ## Whether to report per-cpu stats or not

  percpu = true

  ## Whether to report total system cpu stats or not

  totalcpu = true

  ## If true, collect raw CPU time metrics

  collect_cpu_time = true

  ## If true, compute and report the sum of all non-idle CPU states

  report_active = true


# Read metrics about memory usage

[[inputs.mem]]

  # no configuration


# Read metrics about system load & uptime
```

```
[[inputs.system]]
  ## Uncomment to remove deprecated metrics.
  # fielddrop = ["uptime_format"]

# Read metrics about network interface usage
[[inputs.net]]
  ## By default, telegraf gathers stats from any up interface (excluding loopback)
  ## Setting interfaces will tell it to gather these explicit interfaces,
  ## regardless of status.
  ##
 # interfaces = ["eth0"]

# Read metrics about swap memory usage
[[inputs.swap]]
  # no configuration

# Read metrics about disk usage by mount point
[[inputs.disk]]
  ## By default stats will be gathered for all mount points.
  ## Set mount_points will restrict the stats to only the specified mount points.
  # mount_points = ["/"]

  ## Ignore mount points by filesystem type.
  # ignore_fs = ["tmpfs", "devtmpfs", "devfs", "iso9660", "overlay", "aufs", "squashfs"]

# Get the number of processes and group them by status
[[inputs.processes]]
  # no configuration

# Read metrics about disk IO by device
[[inputs.diskio]]
  ## By default, telegraf will gather stats for all devices including
```

```
  ## disk partitions.
  ## Setting devices will restrict the stats to the specified devices.
  # devices = ["sda", "sdb", "vd*"]
  ## Uncomment the following line if you need disk serial numbers.
  # skip_serial_number = false
```

**Windows Operating System**

To collect Windows OS related metrics and properties as it is in a managed Telegraf agent installation flow and to have localized object types, names, alerts, symptoms, metrics and properties, and so on, the `telegraf.conf` file should be customized and have the following content.

For Windows OS, the `telegraf.conf` file default configurations were changed for Telegraf versions greater than or equal to 1.20.0 and includes Linux related input plugins, like **`inputs.cpu, inputs.disk, inputs.diskio, inputs.kernel, inputs.mem, inputs.processes, inputs.swap, inputs.system`**, and so on. They should be commented and those related to Windows should be uncommented.

```
[[inputs.win_perf_counters]]
PrintValid=true


[[inputs.win_perf_counters.object]]
  ObjectName = "Processor"
  Instances = ["*"]
  Counters = ["% Idle Time", "% Interrupt Time", "% Privileged Time", "% Processor Time",
"% User Time", "Interrupts/sec", "% DPC Time"]
  Measurement = "win.cpu"
  IncludeTotal = true


[[inputs.win_perf_counters.object]]
  ObjectName = "LogicalDisk"
  Instances = ["*"]
  Counters = ["% Disk Read Time", "% Disk Write Time", "% Free Space", "% Idle Time",
"Avg. Disk Bytes/Read", "Avg. Disk Bytes/Write", "Avg. Disk Queue Length", "Avg. Disk sec/
Read", "Avg. Disk sec/Write", "Avg. Disk Write Queue Length", "Avg. Disk Read Queue
Length", "Free Megabytes", "Split IO/Sec"]
  Measurement = "win.disk"


[[inputs.win_perf_counters.object]]
  ObjectName = "Memory"
```

```
  Counters = ["Available Bytes", "Cache Bytes", "Committed Bytes", "Cache Faults/sec",
"Demand Zero Faults/sec", "Page Faults/sec", "Pages/sec", "Transition Faults/sec", "Pool
Nonpaged Bytes", "Pool Paged Bytes"]

  Instances = ["------"]

  Measurement = "win.mem"


[[inputs.win_perf_counters.object]]

  ObjectName = "Network Interface"

  Counters = ["Bytes Received/sec", "Bytes Sent/sec", "Packets Outbound Discarded",
"Packets Outbound Errors", "Packets Received Discarded", "Packets Received Errors",
"Packets Received/sec", "Packets Sent/sec", "Connections Established"]

  Instances = ["*"]

  Measurement = "win.net"

  IncludeTotal = true


[[inputs.win_perf_counters.object]]

  ObjectName = "Paging File"

  Counters = ["% Usage"]

  Instances = ["*"]

  Measurement = "win.paging"

  IncludeTotal = true


  [[inputs.win_perf_counters.object]]

    ObjectName = "Process"

    Counters = ["% Privileged Time", "% Processor Time", "% User Time", "Elapsed Time",
"Handle Count", "IO Read Bytes/sec", "IO Read Operations/sec", "IO Write Bytes/sec", "IO
Write Operations/sec", "Private Bytes", "Thread Count", "Virtual Bytes", "Working Set",
"Working Set - Private"]

    Instances = ["_Total", "telegraf", "w3wp"]         # Replace this with a list of
process names that you want to monitor. "_Total" is all processes combined

    Measurement = "win.process"


[[inputs.win_perf_counters.object]]

  ObjectName = "System"

  Counters = ["Context Switches/sec", "Processes", "Processor Queue Length", "System
Calls/sec", "System Up Time", "Threads"]

  Instances = ["------"]
```

```
  Measurement = "win.system"


[[inputs.win_perf_counters.object]]

  ObjectName = "TCPv4"

  Counters = ["Connection Failures", "Connections Active", "Connections Established",
"Connections Passive", "Connection Reset", "Segments Received/sec", "Segments
Retransmitted/sec", "Segments Sent/sec"]

  Instances = ["------"]

  Measurement = "win.net.tcp"


[[inputs.win_perf_counters.object]]

  ObjectName = "TCPv6"

  Counters = ["Connection Failures", "Connections Active", "Connections Established",
"Connections Passive", "Connection Reset", "Segments Received/sec", "Segments
Retransmitted/sec", "Segments Sent/sec"]

  Instances = ["------"]

  Measurement = "win.net.tcp"


[[inputs.win_perf_counters.object]]

  ObjectName = "UDPv4"

  Counters = ["Datagrams No Port/sec", "Datagrams Received/Errors", "Datagrams Received/
sec", "Datagrams Sent/sec"]

  Instances = ["------"]

  Measurement = "win.net.udp"


[[inputs.win_perf_counters.object]]

  ObjectName = "UDPv6"

  Counters = ["Datagrams No Port/sec", "Datagrams Received/Errors", "Datagrams Received/
sec", "Datagrams Sent/sec"]

  Instances = ["------"]

  Measurement = "win.net.udp"
```

### Sample Configurations

If you do not use the helper script, you must provide specific configurations required on the end point to post metrics to cloud proxy. Based on whether the vCenter VM, the EC2 AWS instance, or the Azure VM is managed or unmanaged, the

following configurations must be provided for open source Telegraf.Based on whether the vCenter VM is managed or unmanaged, the following configurations must be provided for open source Telegraf.You must provide correct values for the variables enclosed in <>.

If you do not use the helper script, you must download mandatory_tags.sh or mandatory_tags.bat and provide the path to the script that you have downloaded and the path to the Telegraf executable. For example, in the following code, you must provide the following: <Path to mandatory_tags.sh/bat> <path to telegraf executable>.

For Linux, download the mandatory_tags.sh from cloud proxy located at `https://<CP_IP>/downloads/salt/mandatory_tags.sh`.

For Windows, download the mandatory_tags.ps1 from cloud proxy located at `https://<CP_IP>/downloads/salt/mandatory_tags.ps1`.

> **NOTE**
> For Windows and Linux, use the relevant cloud proxy IP address/FQDN for <CP_IP> in the preceding commands and location specified. For application monitoring high availability activated collector groups, provide the virtual IP

**Managed vCenter VM**

```
[agent]
  interval = "300s"
  round_interval = true
  metric_batch_size = 1000
  metric_buffer_limit = 2000
  collection_jitter = "0s"
  flush_interval = "60s"
  flush_jitter = "0s"
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<VM_NAME/HOSTNAME>"
  omit_hostname = false


# Configuration for HTTP server to send metrics to
[[outputs.http]]
   url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"
   timeout = "5s"
   method = "POST"
```

```
   insecure_skip_verify = true
   data_format = "wavefront"
   ## Additional HTTP headers
   [outputs.http.headers]
     Content-Type = "text/plain; charset=utf-8"
     vmId = "<VM_MOR>"
     vcid = "<VC_ID>"
     hostname = "<VM_NAME/HOSTNAME>"
     uuid = ""
[[inputs.exec]]
  commands = ["/bin/bash <Path to mandatory_tags.sh/bat> <path to telegraf executable>"]
  timeout = "5s"
  data_format = "influx"
```

**Managed AWS EC2 Instance**

```
[agent]
  interval = "300s"
  round_interval = true
  metric_batch_size = 1000
  metric_buffer_limit = 2000
  collection_jitter = "0s"
  flush_interval = "60s"
  flush_jitter = "0s"
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<EC2_INSTANCE_NAME/HOSTNAME>"
  omit_hostname = true

# Configuration for HTTP server to send metrics to
[[outputs.http]]
```

```
  url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"
  timeout = "5s"
  method = "POST"
  insecure_skip_verify = true
  data_format = "wavefront"
## Additional HTTP headers
[outputs.http.headers]
  Content-Type = "text/plain; charset=utf-8"
  vmId = "i-< EC2_INSTANCE_ID>"
  vcid = "<REAGON>"
  hostname = "<EC2_INSTANCE_NAME/HOSTNAME>"
  uuid = "aws"

[[inputs.exec]]
  commands = ["/bin/bash <Path to mandatory_tags.sh/bat> <path to telegraf executable>"]
  timeout = "5s"
  data_format = "influx"
```

**Managed Azure VM**

```
[agent]
  interval = "300s"
  round_interval = true
  metric_batch_size = 1000
  metric_buffer_limit = 2000
  collection_jitter = "0s"
  flush_interval = "60s"
  flush_jitter = "0s"
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<AZURE_VIRTUAL_MACHINE_NAME/HOSTNAME>"
```

```
  omit_hostname = true


# Configuration for HTTP server to send metrics to

[[outputs.http]]

   url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"

   timeout = "5s"

   method = "POST"

   insecure_skip_verify = true

   data_format = "wavefront"

   ## Additional HTTP headers

   [outputs.http.headers]

     Content-Type = "text/plain; charset=utf-8"

     vmId =
";subscriptions;<SUBSCRIPTION_ID>;resourceGroups;<AZURE_RESOURCE_GROUP>;providers;Microsof
t.Compute;virtualMachines;<AZURE_VIRTUAL_MACHINE_NAME>"

     vcid = "<SUBSCRIPTION_ID>"

     hostname = "<AZURE_VIRTUAL_MACHINE_NAME/HOSTNAME>"

     uuid = "azure"


[[inputs.exec]]

  commands = ["/bin/bash \"/etc/telegraf/telegraf.d/mandatory_tags.sh\" \"/usr/bb

in/telegraf\""]

  timeout = "5s"

  data_format = "influx"
```

**Unmanaged vCenter VMs, AWS EC2 Instances, and Azure VMsUnmanaged vCenter VMs**

```
[agent]

  interval = "300s"

  round_interval = true

  metric_batch_size = 1000

  metric_buffer_limit = 2000

  collection_jitter = "0s"

  flush_interval = "60s"
```

```
  flush_jitter = "0s"
  precision = ""
  debug = false
  quiet = false
  logfile = ""
  hostname = "<VM_NAME/HOSTNAME>"
  omit_hostname = false


# Configuration for HTTP server to send metrics to
[[outputs.http]]
   url = "https://<CP_IP_or_FQDN>:8443/opensource/default/metric"
   timeout = "5s"
   method = "POST"
   insecure_skip_verify = true
   data_format = "wavefront"
   ## Additional HTTP headers
   [outputs.http.headers]
     Content-Type = "text/plain; charset=utf-8"
     uuid = "<UUID>"
     ip = "<IP_ADDRESS>"
     hostname = "<VM_NAME/HOSTNAME>"
[[inputs.exec]]
  commands = ["/bin/bash <Path to mandatory_tags.sh/bat> <path to telegraf executable>"]
  timeout = "5s"
  data_format = "influx"
```

### Telegraf Configuration Details for Supported Application Services

For application services and custom monitoring plugins supported by VMware Aria OperationsVMware Cloud Foundation Operations, there are configuration details that you must follow.

**Table 92: Configuration Details for Supported Application Services**

| Configuration Details for Supported Application Services |
| --- |
| **Active Directory**<br>Here are the configuration details:<br>`[[inputs.win_perf_counters]]`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`    ObjectName = "DirectoryServices"`<br>`    Counters = ["LDAP Active Threads","LDAP Client Sessions","LDAP Writes/sec","LDAP Searches/sec","LDAP Successful Binds/sec","LDAP New Connections/sec","LDAP Closed Connections/sec","LDAP UDP operations/sec","DS Threads in Use","DS Directory Writes/sec","DS Directory Reads/sec","DS Directory Searches/sec","DS Client Binds/sec","DS Server Binds/sec","DRA Pending Replication Synchronizations","DRA Sync Requests Made","DRA Sync Requests Successful","DRA Pending Replication Operations","DRA Inbound Objects/sec","DRA Inbound Bytes Total/sec","DRA Outbound Objects/sec","DRA Outbound Bytes Total/sec","Base Searches/sec","Database adds/sec","Database deletes/sec","Database modifys/sec","Database recycles/sec"]`<br>`    Instances = ["*"]`<br>`    Measurement = "ad.active.directory"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`    ObjectName = "Security System-Wide Statistics"`<br>`    Counters = ["NTLM Authentications","Kerberos Authentications","KDC AS Requests","KDC TGS Requests","Digest Authentications"]`<br>`    Instances = ["------"]`<br>`    Measurement = "ad.security.statistics"`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`    ObjectName = "Database"`<br>`    Instances = ["*"]`<br>`    Counters = ["Database Cache % Hit","Database Cache Page Faults/sec","Database Page Faults/sec","Database Cache Page Fault Stalls/sec","Database Cache Size"]`<br>`    Measurement = "ad.database"`<br>`    #IncludeTotal=false #Set to true to include _Total instance when querying for all (*).`<br><br>`  [[inputs.win_perf_counters.object]]`<br>`    ObjectName = "DFS Namespace Service Referrals"` |

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
    Instances = ["*"]

    Counters = ["Requests Processed","Requests Failed","Avg. Response Time"]

    Measurement = "ad.dfsn"


  [[inputs.win_perf_counters.object]]

    ObjectName = "DFS Replication Service Volumes"

    Instances = ["*"]

    Counters = ["Data Lookups","Database Lookups","Database Commits"]

    Measurement = "ad.dfs.replications"


  [[inputs.win_perf_counters.object]]

    ObjectName = "DNS"

    Counters = ["Dynamic Update Received","Dynamic Update Rejected","Recursive
Queries","Recursive Queries Failure","Recursive Query Failure","Secure Update
Failure","Total Query Received","Total Response Sent"]

    Instances = ["*"]

    Measurement = "ad.dns"
```

**ActiveMQ**

Here are the configuration details:

```
# Read Apache ActiveMQ information.

[[inputs.jolokia2_agent]]

  urls = ["http://localhost:8161/api/jolokia"]

  name_prefix = "activemq."

  plugin_name_override="activemq"

     username = "user"

     password = "user"


### JVM Generic


  [[inputs.jolokia2_agent.metric]]

    name = "OperatingSystem"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  mbean = "java.lang:type=OperatingSystem"
  paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]


[[inputs.jolokia2_agent.metric]]
  name = "OperatingSystem"
  mbean = "java.lang:type=OperatingSystem"
  paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]


[[inputs.jolokia2_agent.metric]]
  name  = "jvm_runtime"
  mbean = "java.lang:type=Runtime"
  paths = ["Uptime"]


[[inputs.jolokia2_agent.metric]]
  name  = "jvm_memory"
  mbean = "java.lang:type=Memory"
  paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]


[[inputs.jolokia2_agent.metric]]
  name     = "jvm_garbage_collector"
  mbean    = "java.lang:name=*,type=GarbageCollector"
  paths    = ["CollectionTime", "CollectionCount"]
  tag_keys = ["name"]


[[inputs.jolokia2_agent.metric]]
  name       = "jvm_memory_pool"
  mbean      = "java.lang:name=*,type=MemoryPool"
  paths      = ["Usage", "PeakUsage", "CollectionUsage"]
  tag_keys   = ["name"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
    tag_prefix = "pool_"


### ACTIVEMQ

  [[inputs.jolokia2_agent.metric]]

    name     = "queue"

    mbean    =
"org.apache.activemq:brokerName=*,destinationName=*,destinationType=Queue,type=Broker"

    paths    =
["QueueSize","EnqueueCount","ConsumerCount","DispatchCount","DequeueCount","ProducerCoun
t","InFlightCount"]

    tag_keys = ["brokerName","destinationName"]



  [[inputs.jolokia2_agent.metric]]

    name     = "topic"

    mbean    =
"org.apache.activemq:brokerName=*,destinationName=*,destinationType=Topic,type=Broker"

    paths    =
["ProducerCount","DequeueCount","ConsumerCount","QueueSize","EnqueueCount"]

    tag_keys = ["brokerName","destinationName"]



  [[inputs.jolokia2_agent.metric]]

    name     = "broker"

    mbean    = "org.apache.activemq:brokerName=*,type=Broker"

    paths    =
["TotalConsumerCount","TotalMessageCount","TotalEnqueueCount","TotalDequeueCount","Memor
yLimit","MemoryPercentUsage","StoreLimit","StorePercentUsage","TempPercentUsage","TempLi
mit"]

    tag_keys = ["brokerName"]



  [[inputs.jolokia2_agent.metric]]

    name  = "java_class_loading"

    mbean = "java.lang:type=ClassLoading"

    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths = ["ThreadCount"]

  [[inputs.jolokia2_agent.metric]]
    name  = "java_buffer_pool"
    mbean = "java.nio:name=*,type=BufferPool"
    paths = ["Count", "MemoryUsed", "TotalCapacity"]
    tag_keys = ["name"]
```

**Apache HTTPD**

Here are the configuration details:

```
# Read Apache status information (mod_status)
[[inputs.apache]]

  ## An array of URLs to gather from, must be directed at the machine
  ## readable version of the mod_status page including the auto query string.
  ## Default is "http://localhost/server-status?auto".
  ##
    ##urls = ["http://localhost/server-status?auto"]
  ##
  urls = ["http://127.0.0.1:80/server-status?auto"]
  ## Credentials for basic HTTP authentication.
    # username = "myuser"
    #password = "mypassword"

  ## Maximum time to receive response.
  # response_timeout = "5s"
```

*Table continued on next page*

*Continued from previous page*

| Configuration Details for Supported Application Services |
| --- |
| ```<br>## Optional SSL Config<br>    # ssl_ca = "/etc/telegraf/ca.pem"<br>  # ssl_cert = "/etc/telegraf/cert.pem"<br>  # ssl_key = "/etc/telegraf/key.pem"<br><br>## Use SSL but skip chain & host verification<br>  insecure_skip_verify = true<br>``` |
| **Cassandra**<br>Here are the configuration details:<br>```<br># Read Cassandra metrics through Jolokia<br>[[inputs.cassandra]]<br>  context = "/jolokia/read"<br>  ## List of cassandra servers exposing jolokia read service<br>  servers = ["myuser:mypassword@10.10.10.1:8778","10.10.10.2:8778",":8778"]<br>  ## List of metrics collected on above servers<br>  ## Each metric consists of a jmx path.<br>  ## This will collect all heap memory usage metrics from the jvm and<br>  ## ReadLatency metrics for all keyspaces and tables.<br>  ## "type=Table" in the query works with Cassandra3.0. Older versions might<br>  ## need to use "type=ColumnFamily"<br>  metrics  = [<br>  "/java.lang:type=Memory/HeapMemoryUsage",<br>  "/org.apache.cassandra.metrics:type=Table,keyspace=*,scope=*,name=ReadLatency"<br>  ]<br>``` |
| **Hyper-V**<br>Here are the configuration details:<br>```<br>[[inputs.win_perf_counters]]<br><br>  [[inputs.win_perf_counters.object]]<br>  ObjectName = "Hyper-V Virtual Machine Health Summary"<br>  Instances = ["------"]<br>``` |

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
      Measurement = "hyperv.vm.health"
      Counters = ["Health Ok", "Health Critical"]


    [[inputs.win_perf_counters.object]]
      ObjectName = "Hyper-V Dynamic Memory VM"
      Instances = ["*"]
      Measurement = "hyperv.vm.memory"
      Counters = ["Physical Memory", "Added Memory", "Guest Visible Physical Memory"]


    [[inputs.win_perf_counters.object]]
      ObjectName = "Hyper-V Hypervisor Virtual Processor"
      Instances = ["*"]
      Measurement = "hyperv.hypervisor.virtual.processor"
      Counters = ["% Guest Run Time", "% Hypervisor Run Time", "% Total Run Time"]


    [[inputs.win_perf_counters.object]]
      ObjectName = "Hyper-V Hypervisor Root Virtual Processor"
      Instances = ["*"]
      Measurement = "hyperv.hypervisor.root.virtual.processor"
      Counters = ["% Guest Run Time", "% Hypervisor Run Time", "% Total Run Time"]
      IncludeTotal = true


    [[inputs.win_perf_counters.object]]
      ObjectName = "Hyper-V Virtual IDE Controller (Emulated)"
      Instances = ["*"]
      Measurement = "hyperv.virtual.ide.controller"
      Counters = ["Write Bytes/sec", "Read Bytes/sec", "Written Sectors/sec", "Read
Sectors/sec"]


    [[inputs.win_perf_counters.object]]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
       ObjectName = "Hyper-V Virtual Network Adapter"

       Instances = ["*"]

       Measurement = "hyperv.virtual.net.adapter"

       Counters = ["Bytes/sec", "Bytes Received/sec", "Bytes Sent/Sec", "Packets Sent/
sec", "Packets Received/sec", "Packets/sec"]


     [[inputs.win_perf_counters.object]]

       ObjectName = "Hyper-V Legacy Network Adapter"

       Instances = ["*"]

       Measurement = "hyperv.legacy.net.adapter"

       Counters = ["Bytes Dropped", "Bytes Received/sec", "Bytes Sent/Sec"]


     [[inputs.win_perf_counters.object]]

       ObjectName = "Hyper-V Hypervisor Partition"

       Instances = ["*"]

       Measurement = "hyperv.hypervisor.partition"

       Counters = ["Virtual Processors"]


     [[inputs.win_perf_counters.object]]

       ObjectName = "Hyper-V Virtual Storage Device"

       Instances = ["*"]

       Measurement = "hyperv.virtual.storage.device"

       Counters = [

                  "Maximum Bandwidth", "Read Bytes/sec", "Write Bytes/sec", "Queue
Length",

                  "Lower Latency", "Minimum IO Rate", "Maximum IO Rate", "Latency",
"Throughput",

                  "Lower Queue Length", "Queue Length", "Normalized Throughput", "Write
Operations/Sec",

                  "Read Operations/Sec", "Write Bytes/sec", "Read Bytes/sec", "Error
Count",

                  "Flush Count", "Write Count", "Read Count"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
               ]


     [[inputs.win_perf_counters.object]]

       ObjectName = "Processor"

       Instances = ["*"]

       Counters = ["% Idle Time", "% Interrupt Time", "% Privileged Time", "% Processor
Time", "% User Time", "Interrupts/sec"]

       Measurement = "hyperv.host.cpu"

       IncludeTotal = true


     [[inputs.win_perf_counters.object]]

       ObjectName = "LogicalDisk"

       Instances = ["*"]

       Counters = ["% Disk Read Time", "% Disk Write Time", "% Free Space", "% Idle
Time", "Avg. Disk Bytes/Read", "Avg. Disk Bytes/Write", "Avg. Disk Queue Length", "Avg.
Disk sec/Read", "Avg. Disk sec/Write", "Avg. Disk Write Queue Length", "Free Megabytes",
"Split IO/Sec"]

       Measurement = "hyperv.host.disk"

       IncludeTotal = true


     [[inputs.win_perf_counters.object]]

       ObjectName = "Memory"

       Counters = ["Available Bytes", "Cache Bytes", "Committed Bytes", "Cache Faults/
sec", "Demand Zero Faults/sec", "Page Faults/sec", "Pages/sec", "Transition Faults/sec",
"Pool Nonpaged Bytes", "Pool Paged Bytes"]

       Instances = ["------"]

       Measurement = "hyperv.host.mem"


     [[inputs.win_perf_counters.object]]

       ObjectName = "Network Interface"

       Counters = ["Bytes Received/sec", "Bytes Sent/sec", "Packets Outbound Discarded",
"Packets Outbound Errors", "Packets Received Discarded", "Packets Received Errors",
"Packets Received/sec", "Packets Sent/sec", "Bytes Total/sec", "Current Bandwidth",
"Output Queue Length"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
       Instances = ["*"]

       Measurement = "hyperv.host.net"

       IncludeTotal = true


     [[inputs.win_perf_counters.object]]

       ObjectName = "System"

       Counters = ["Context Switches/sec", "Processes", "Processor Queue Length",
"System Calls/sec", "System Up Time", "Threads"]

       Instances = ["------"]

       Measurement = "hyperv.host.system"


     [[inputs.win_perf_counters.object]]

       ObjectName = "Process"

       Counters = ["% Privileged Time", "% Processor Time", "% User Time", "Elapsed
Time", "Handle Count", "IO Read Bytes/sec", "IO Read Operations/sec", "IO Write Bytes/
sec", "IO Write Operations/sec", "Private Bytes", "Thread Count", "Virtual Bytes",
"Working Set", "Working Set - Private"]

       Instances = ["_Total"]

       Measurement = "hyperv.host.process"
```

**Java Plugin**

Here are the configuration details:

```
[[inputs.jolokia2_agent]]

  # Prefix to attach to the measurement name

  name_prefix = "java."

  # Add agents URLs to query

  urls = ["http://localhost:8080/jolokia"]

  #username and password are mandatory for Jolokia 1.6 or later

  #username = <jolokia role username>

  #password = <jolokia role password>

  # response_timeout = "5s"

  ## Optional TLS config

  # tls_ca   = "/var/private/ca.pem"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  # tls_cert = "/var/private/client.pem"
  # tls_key  = "/var/private/client-key.pem"
  # insecure_skip_verify = false
  ### JVM Generic
  [[inputs.jolokia2_agent.metric]]
    name  = "OperatingSystem"
    mbean = "java.lang:type=OperatingSystem"
    paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]
  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_runtime"
    mbean = "java.lang:type=Runtime"
    paths = ["Uptime"]
  [[inputs.jolokia2_agent.metric]]
    name  = "jvm_memory"
    mbean = "java.lang:type=Memory"
    paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]
  [[inputs.jolokia2_agent.metric]]
    name     = "jvm_garbage_collector"
    mbean    = "java.lang:name=*,type=GarbageCollector"
    paths    = ["CollectionTime", "CollectionCount"]
    tag_keys = ["name"]
  [[inputs.jolokia2_agent.metric]]
    name       = "jvm_memory_pool"
    mbean      = "java.lang:name=*,type=MemoryPool"
    paths      = ["Usage", "PeakUsage", "CollectionUsage"]
    tag_keys   = ["name"]
    tag_prefix = "pool_"
  ### TOMCAT
  [[inputs.jolokia2_agent.metric]]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
    name     = "GlobalRequestProcessor"

    mbean    = "Catalina:name=*,type=GlobalRequestProcessor"

    paths    =
["requestCount","bytesReceived","bytesSent","processingTime","errorCount"]

    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]

    name     = "JspMonitor"

    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,name=jsp,type=JspMonitor"

    paths    = ["jspReloadCount","jspCount","jspUnloadCount"]

    tag_keys = ["J2EEApplication","J2EEServer","WebModule"]

  [[inputs.jolokia2_agent.metric]]

    name     = "ThreadPool"

    mbean    = "Catalina:name=*,type=ThreadPool"

    paths    = ["maxThreads","currentThreadCount","currentThreadsBusy"]

    tag_keys = ["name"]

  [[inputs.jolokia2_agent.metric]]

    name     = "Servlet"

    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,j2eeType=Servlet,name=*"

    paths    = ["processingTime","errorCount","requestCount"]

    tag_keys = ["name","J2EEApplication","J2EEServer","WebModule"]

  [[inputs.jolokia2_agent.metric]]

    name     = "Cache"

    mbean    = "Catalina:context=*,host=*,name=Cache,type=WebResourceRoot"

    paths    = ["hitCount","lookupCount"]

    tag_keys = ["context","host"]
```

**JBoss Server**

Here are the configuration details:

```
# Read JBoss status information (mod_status)

[[inputs.jolokia2_agent]]

  ##urls = ["http://localhost:8380/jolokia"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
urls = ["http://127.0.0.1:8082/jolokia"]
name_prefix = "jboss."

## Optional SSL Config
  # ssl_ca = "/etc/telegraf/ca.pem"
  # ssl_cert = "/etc/telegraf/cert.pem"
  # ssl_key = "/etc/telegraf/key.pem"
## Use SSL but skip chain & host verification
   insecure_skip_verify = true

### JVM Generic

[[inputs.jolokia2_agent.metric]]
  name = "OperatingSystem"
  mbean = "java.lang:type=OperatingSystem"
  paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]

[[inputs.jolokia2_agent.metric]]
  name = "OperatingSystem"
  mbean = "java.lang:type=OperatingSystem"
  paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]

[[inputs.jolokia2_agent.metric]]
  name  = "jvm_runtime"
  mbean = "java.lang:type=Runtime"
  paths = ["Uptime"]

[[inputs.jolokia2_agent.metric]]
  name  = "jvm_memory"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  mbean = "java.lang:type=Memory"
  paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]


[[inputs.jolokia2_agent.metric]]
  name     = "jvm_garbage_collector"
  mbean    = "java.lang:name=*,type=GarbageCollector"
  paths    = ["CollectionTime", "CollectionCount"]
  tag_keys = ["name"]


[[inputs.jolokia2_agent.metric]]
  name       = "jvm_memory_pool"
  mbean      = "java.lang:name=*,type=MemoryPool"
  paths      = ["Usage", "PeakUsage", "CollectionUsage"]
  tag_keys   = ["name"]
  tag_prefix = "pool_"


### JBOSS


[[inputs.jolokia2_agent.metric]]
  name     = "connectors.http"
  mbean    = "jboss.as:https-listener=*,server=*,subsystem=undertow"
  paths    = ["bytesReceived","bytesSent","errorCount","requestCount"]
  tag_keys = ["server","https-listener"]


[[inputs.jolokia2_agent.metric]]
  name     = "connectors.http"
  mbean    = "jboss.as:http-listener=*,server=*,subsystem=undertow"
  paths    = ["bytesReceived","bytesSent","errorCount","requestCount"]
  tag_keys = ["server","http-listener"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  [[inputs.jolokia2_agent.metric]]
    name     = "datasource.jdbc"
    mbean    = "jboss.as:data-source=*,statistics=jdbc,subsystem=datasources"
    paths    =
["PreparedStatementCacheAccessCount","PreparedStatementCacheHitCount","PreparedStatement
CacheMissCount"]
    tag_keys = ["data-source"]


  [[inputs.jolokia2_agent.metric]]
    name     = "datasource.pool"
    mbean    = "jboss.as:data-source=*,statistics=pool,subsystem=datasources"
    paths    = ["AvailableCount","ActiveCount","MaxUsedCount"]
    tag_keys = ["data-source"]


  [[inputs.jolokia2_agent.metric]]
    name = "thread.count"
    mbean  = "jboss.as:subsystem=*,thread-pool=*"
    paths =
["completedTaskCount","currentThreadCount","maxThreads","keepaliveTime","largestThreadCo
unt","activeCount","taskCount","rejectedCount"]


  [[inputs.jolokia2_agent.metric]]
    name  = "java_class_loading"
    mbean = "java.lang:type=ClassLoading"
    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]


  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths =
["ThreadCount","TotalStartedThreadCount","DaemonThreadCount","PeakThreadCount"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  [[inputs.jolokia2_agent.metric]]
    name  = "java_buffer_pool"
    mbean = "java.nio:name=*,type=BufferPool"
    paths = ["Count", "MemoryUsed", "TotalCapacity"]
    tag_keys = ["name"]
```

**Microsoft IIS**

Here are the configuration details:

```
# Read MSIIS status information (mod_status)
[[inputs.win_perf_counters]]


     [[inputs.win_perf_counters.object]]
       # IIS, Web Service
       ObjectName = "Web Service"
       Counters = [
                     "Service Uptime", "Current Connections", "Bytes Sent/sec", "Total
Bytes Sent", "Connection Attempts/sec",
                     "Bytes Received/sec", "Total Bytes Received", "Bytes Total/sec",
"Total Bytes Transferred",
                     "Get Requests/sec","Total Get Requests", "Post Requests/sec","Total
Post Requests",
                     "Put Requests/sec","Total Put Requests", "Delete Requests/
sec","Total Delete Requests",
                     "Anonymous Users/sec", "NonAnonymous Users/sec", "Files Sent/sec",
"Total Files Sent",
                     "Files Received/sec", "Total Files Received", "Files/sec", "Total
Files Transferred",
                     "Not Found Errors/sec", "Locked Errors/sec", "Total Method Requests/
sec"
                   ]
       Instances = ["*"]
       Measurement = "iis.websvc"
       IncludeTotal=true #Set to false to not include _Total instance.
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
     [[inputs.win_perf_counters.object]]

       # Web Service Cache / IIS

       ObjectName = "Web Service Cache"

       Counters = [

                     "Current Files Cached", "Active Flushed Entries", "Total Files
Cached", "Total Flushed Files",

                     "File Cache Hits", "File Cache Misses", "File Cache Hits %", "File
Cache Flushes",

                     "Current File Cache Memory Usage", "Maximum File Cache Memory
Usage", "Current URIs Cached",

                     "Total URIs Cached", "Total Flushed URIs", "URI Cache Hits", "URI
Cache Misses", "URI Cache Hits %",

                     "URI Cache Flushes", "Current Metadata Cached", "Total Metadata
Cached", "Total Flushed Metadata",

                     "Metadata Cache Hits", "Metadata Cache Misses", "Metadata Cache Hits
%", "Metadata Cache Flushes",

                     "Kernel: Current URIs Cached", "Kernel: Total URIs Cached", "Kernel:
Total Flushed URIs",

                     "Kernel: URI Cache Hits", "Kernel: Uri Cache Hits/sec", "Kernel: URI
Cache Misses", "Kernel: URI Cache Hits %",

                     "Kernel: URI Cache Flushes", "Cache Memory Usage"

                   ]

       Instances = ["*"]

       Measurement = "iis.websvc.cache"

       IncludeTotal=true #Set to false to not include _Total instance.


     [[inputs.win_perf_counters.object]]

       # IIS, ASP.NET

       ObjectName = "ASP.NET"

       Counters = ["Application Restarts","Request Wait Time","Requests
Current","Requests Queued","Requests Rejected"]

       Instances = ["*"]

       Measurement = "iis.aspnet"

       #IncludeTotal=false #Set to true to include _Total instance when querying for all
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
(*).



     [[inputs.win_perf_counters.object]]

       # IIS, ASP.NET Applications

       ObjectName = "ASP.NET Applications"

       Counters = ["Cache Total Entries","Cache Total Hit Ratio","Output Cache
Entries","Errors Total/Sec","Requests Executing","Requests in Application
Queue","Requests/Sec"]

       Instances = ["*"]

       Measurement = "iis.aspnet.app"

       #IncludeTotal=false #Set to true to include _Total instance when querying for all
(*).



     [[inputs.win_perf_counters.object]]

       ObjectName = ".NET CLR Exceptions"

       Counters = ["# of Exceps Thrown", "# of Exceps Thrown / Sec", "# of Filters /
Sec", "# of Finallys / Sec", "Throw to Catch Depth / Sec"]

       Instances = ["*"]

       Measurement = "iis.dotnet.exception"



     [[inputs.win_perf_counters.object]]

       ObjectName = ".NET CLR Jit"

       Counters = ["% Time in Jit","IL Bytes Jitted / sec"]

       Instances = ["*"]

       Measurement = "iis.dotnet.jit"



     [[inputs.win_perf_counters.object]]

       ObjectName = ".NET CLR Loading"

       Counters = ["% Time Loading"]

       Instances = ["*"]

       Measurement = "iis.dotnet.loading"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
     [[inputs.win_perf_counters.object]]

       ObjectName = ".NET CLR LocksAndThreads"

       Counters = ["# of current logical Threads","# of current physical Threads","# of
current recognized threads","# of total recognized threads","Queue Length / sec","Total
# of Contentions","Current Queue Length"]

       Instances = ["*"]

       Measurement = "iis.dotnet.lock"



     [[inputs.win_perf_counters.object]]

       ObjectName = ".NET CLR Memory"

       Counters = ["% Time in GC","# Bytes in all Heaps","# Gen 0 Collections","# Gen 1
Collections","# Gen 2 Collections","# Induced GC","Allocated Bytes/sec","Finalization
Survivors","Gen 0 heap size","Gen 1 heap size","Gen 2 heap size","Large Object Heap
size","# of Pinned Objects"]

       Instances = ["*"]

       Measurement = "iis.dotnet.memory"



     [[inputs.win_perf_counters.object]]

       ObjectName = ".NET CLR Security"

       Counters = ["% Time in RT checks","Stack Walk Depth","Total Runtime Checks"]

       Instances = ["*"]

       Measurement = "iis.dotnet.security"



     [[inputs.win_perf_counters.object]]

       # HTTP Service request queues in the Kernel before being handed over to User
Mode.

       ObjectName = "HTTP Service Request Queues"

       Instances = ["*"]

       Counters = ["CurrentQueueSize","RejectedRequests"]

       Measurement = "iis.http.queues"

       #IncludeTotal=false #Set to true to include _Total instance when querying for all
(*).
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
     [[inputs.win_perf_counters.object]]

       # Process metrics, in this case for IIS only

       ObjectName = "Process"

       Counters = ["% Processor Time","Handle Count","Private Bytes","Thread
Count","Virtual Bytes","Working Set"]

       Instances = ["w3wp"]

       Measurement = "iis.win.proc"

       #IncludeTotal=false #Set to true to include _Total instance when querying for all
(*).
```

**Microsoft SQL Server**

Here are the configuration details:

```
[[inputs.sqlserver]]

   name_prefix = "MSSQL."



       #servers = [

       #"Server=<servername>;Port=1433;User Id=telegraf;Password=<mystrongpassword from
step 2>;app name=telegraf;log=1;",

       #"Server=<servername>;Port=1433;User Id=telegraf;Password=<mystrongpassword from
step 2>;app name=telegraf;log=1;"

      #]

       servers = [

        "Server=localhost;Port=1433;User Id=sa;Password=Password;app
name=telegraf;log=1;"

         ]



   namepass = ["Rows*writes*bytes*sec*", "Rows*reads*bytes*sec*", "Rows*writes*sec*",
"Rows*reads*sec*", "Query*User*counter*", "Buffer*cache*hit*ratio*",
"Page*life*expectancy*", "Page*lookups*sec*", "Page*reads*sec*", "Page*writes*sec*",
"Lazy*writes*sec*", "Checkpoint*pages*sec*", "Log*Apply*Ready*Queue*",
"Data*File*s*Size*KB*", "Log*File*s*", "XTP*Memory*Used*KB*", "Log*Flushes*sec*",
"Write*Transactions*sec*", "Transactions*sec*", "Log*Flush*Wait*Time*",
"Active*Transactions*", "Log*Bytes*Flushed*sec*", "Processes*blocked*",
"User*Connections*", "Logins*sec*", "Logouts*sec*", "Active*Temp*Tables*",
"Temp*Tables*Creation*Rate*", "Batch*Requests*sec*", "SQL*Compilations*sec*",
"SQL*Re*Compilations*sec", "Stored*Procedures*Invoked*sec*", "Target*Server*Memory*KB*",
"Total*Server*Memory*KB*", "SQL*Cache*Memory*KB*", "Log*Pool*Memory*KB*",
"Connection*Memory*KB*", "Lock*Memory*KB*", "Memory*Grants*Pending*",
"Active*memory*grant*amount*KB*", "Disk*Read*Bytes*sec*", "Disk*Read*IO*Throttled*sec*",
"Disk*Read*IO*sec*", "Disk*Write*Bytes*sec*", "Disk*Write*IO*Throttled*sec*",
"Used*memory*KB*", "CPU*usage*", "Free*Space*in*tempdb*KB*", "Version*Store*Size*KB*",
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
"Transactions*", "Blocked*tasks*", "Active*requests*", "Queued*requests*",
"Requests*completed*sec*", "Number*of*Deadlocks*sec*", "Lock*Wait*Time*ms*",
"Lock*Waits*sec*", "Lock*Requests*sec*", "Average*Wait*Time*ms*", "Index*Searches*sec*",
"Page*Splits*sec*", "Full*Scans*sec*", "CPU*", "Wait*time*ms*", "Wait*tasks*", "State*",
"Recovery*Model*"]
```

**MongoDB**

Here are the configuration details:

```
# Read metrics from one or many MongoDB servers

[[inputs.mongodb]]

  ## An array of URI to gather stats about. Specify an ip or hostname

  ## with optional port add password. ie,

  ##   mongodb://user:auth_key@10.10.3.30:27017,

  ##   mongodb://10.10.3.33:18832,

  ##   10.0.0.1:10000, etc.

  ##servers = ["127.0.0.1:27017"]


  servers = ["localhost:27017"]


  gather_perdb_stats = true


  ## Optional SSL Config

    # ssl_ca = "/etc/telegraf/ca.pem"

    # ssl_cert = "/etc/telegraf/cert.pem"

    # ssl_key = "/etc/telegraf/key.pem"


  ## Use SSL but skip chain & host verification

    # insecure_skip_verify = false
```

**MS Exchange Server**

Here are the configuration details:

```
# Read MS Exchange status information (mod_status)

    [[inputs.win_services]]

      name_prefix = "msexchange."

      service_names = [
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
        "MSExchangeTransport",
        "MSExchangeTransportLogSearch",
        "MSExchangeUM",
        "MSExchangeUMCR",
        "MSExchangeThrottling",
        "MSExchangeServiceHost",
        "MSExchangeFastSearch",
        "MSExchangeRepl",
        "MSExchangeRPC",
        "MSExchangePop3",
        "MSExchangeSubmission",
        "MSExchangeDelivery",
        "MSExchangeMailboxReplication",
        "MSExchangeMailboxAssistants",
        "MSExchangeIS",
        "MSExchangeImap4",
        "MSExchangeHM",
        "MSExchangeFrontendTransport",
        "MSExchangeEdgeSync",
        "MSExchangeDiagnostics",
        "MSExchangeAntispamUpdate",
        "MSExchangeADTopology"
    ]



  [[inputs.win_perf_counters]]
    plugin_name_override="msexchange"
    [[inputs.win_perf_counters.object]]
      ObjectName = "Process"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
        Counters = ["% Processor Time","Handle Count","Private Bytes","Thread
Count","Virtual Bytes","Working Set"]

        Instances = [

          "MSExchangeTransport",

          "MSExchangeTransportLogSearch",

          "umservice",

          "Microsoft.Exchange.UM.CallRouter",

          "MSExchangeThrottling",

          "Microsoft.Exchange.ServiceHost",

          "Microsoft.Exchange.Search.Service",

          "msexchangerepl",

          "Microsoft.Exchange.RpcClientAccess.Service",

          "Microsoft.Exchange.Pop3Service",

          "MSExchangeSubmission",

          "MSExchangeDelivery",

          "MSExchangeMailboxReplication",

          "MSExchangeMailboxAssistants",

          "Microsoft.Exchange.Store.Service",

          "Microsoft.Exchange.Imap4Service",

          "MSExchangeHMHost",

          "MSExchangeFrontendTransport",

          "Microsoft.Exchange.EdgeSyncSvc",

          "Microsoft.Exchange.Diagnostics.Service",

          "Microsoft.Exchange.AntispamUpdateSvc",

          "Microsoft.Exchange.Directory.TopologyService"

        ]

        Measurement = "msexchange.process"


      [[inputs.win_perf_counters.object]]

        ObjectName = "Database"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
        Counters = ["I/O Log Writes Average Latency","Log Threads Waiting","Log Record
Stalls/sec",
                "Database Cache % Hit","Database Page Fault Stalls/sec",
                    "I/O Database Writes Average Latency","I/O Database Reads Average
Latency"]
        Instances = ["*"]
        Measurement = "msexchange"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Database"
        Counters = [
          "I/O Log Writes Average Latency",
          "I/O Log Reads Average Latency",
          "Log Threads Waiting",
          "Log Record Stalls/sec",
          "Database Cache % Hit",
          "Database Page Fault Stalls/sec",
          "I/O Database Writes Average Latency",
          "I/O Database Reads Average Latency"
        ]
        Instances = ["*"]
        Measurement = "msexchange.database"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Database ==> Instances"
        Counters = [
          "Log Threads Waiting",
          "Log Record Stalls/sec",
          "I/O Database Writes Average Latency",
          "I/O Database Reads Average Latency"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
        ]

        Instances = ["*"]

        Measurement = "msexchange.databaseinstance"


      [[inputs.win_perf_counters.object]]

        ObjectName = "MSExchange ADAccess Domain Controllers"

        Counters = ["LDAP Read Time","LDAP Search Time","LDAP Searches Timed Out per
Minute","Long Running LDAP Operations/min"]

        Instances = ["*"]

        Measurement = "msexchange.adaccessdomaincontrollers"


      [[inputs.win_perf_counters.object]]

        ObjectName = "MSExchangeUMAvailability"

        Counters = ["Directory Access Failures","Total Inbound Calls Rejected by the UM
Service",

                "Total Inbound Calls Rejected by the UM Worker Process"]

        Instances = ["*"]

        Measurement = "msexchange.umavailability"


      [[inputs.win_perf_counters.object]]

        ObjectName = "MSExchange Availability Service"

        Counters = ["Average Time to Process a Free Busy Request"]

        Instances = ["*"]

        Measurement = "msexchange.availabilityservice"


      [[inputs.win_perf_counters.object]]

        ObjectName = "MSExchange OWA"

        Counters = ["Current Unique Users", "Requests/sec", "Average Search Time"]

        Instances = ["*"]

        Measurement = "msexchange.owa"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeFrontEndTransport SmtpSend"
        Counters = ["Bytes Sent Total", "Messages Sent Total", "Message Bytes Sent
Total"]
        Instances = ["*"]
        Measurement = "msexchange.frontendtransportsmtpsend"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeTransport SmtpSend"
        Counters = ["Message Bytes Sent Total", "Messages Sent Total"]
        Instances = ["_total"]
        Measurement = "msexchange.transportsmtpsend"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeTransport Queues"
        Counters = ["Active Mailbox Delivery Queue Length", "Retry Mailbox Delivery
Queue Length", "Submission Queue Length"]
        Instances = ["_total"]
        Measurement = "msexchange.transportqueues"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeIS Store"
        Counters = ["RPC Average Latency", "RPC Requests","RPC Operations/sec", ]
        Instances = ["*"]
        Measurement = "msexchange.isstore"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeIS Client Type"
        Counters = ["RPC Average Latency", "RPC Operations/sec" ]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
        Instances = ["*"]
        Measurement = "msexchange.isclient"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Store Interface"
        Counters = ["RPC Latency average (msec)"]
        Instances = ["_total"]
        Measurement = "msexchange.isclientinterface"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange ADAccess Processes"
        Counters = ["LDAP Read Time", "LDAP Search Time"]
        Instances = ["*"]
        Measurement = "msexchange.adaccessprocesses"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange ActiveSync"
        Counters = ["Mailbox Search Total", "Requests/sec", "Ping Commands Pending",
"Current Requests",
                    "Average Request Time", "Sync Commands/sec"]
        Instances = ["*"]
        Measurement = "msexchange.activesync"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange Active Manager Server"
        Counters = ["Server-Side Calls/sec", "Active Manager Database State writes to
Persistent storage/sec",
                    "GetServerForDatabase Server-Side Calls", "Total Number of
Databases", "Active Manager Role"]
        Instances = ["*"]
        Measurement = "msexchange.activemanagerserver"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeAutodiscover"
        Counters = ["Requests/sec"]
        Instances = ["*"]
        Measurement = "msexchange.autodiscover"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchangeWS"
        Counters = ["Requests/sec"]
        Instances = ["*"]
        Measurement = "msexchange.ws"


      [[inputs.win_perf_counters.object]]
        ObjectName = "Web Service"
        Counters = ["Current Connections","Connection Attempts/sec","Other Request
Methods/sec"]
        Instances = ["*"]
        Measurement = "msexchange.ws"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange WorkloadManagement Workloads"
        Counters = ["ActiveTasks","CompletedTasks","QueuedTasks"]
        Instances = ["*"]
        Measurement = "msexchange.workload"


      [[inputs.win_perf_counters.object]]
        ObjectName = "ASP.NET"
        Counters = ["Application Restarts","Worker Process Restarts","Request Wait
Time"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
        Instances = ["*"]
        Measurement = "msexchange.aspnet"


      [[inputs.win_perf_counters.object]]
        ObjectName = "MSExchange RpcClientAccess"
        Counters = ["RPC Averaged Latency", "RPC Requests", "Active User Count",
"Connection Count",
                    "RPC Operations/sec", "User Count"]
        Instances = ["*"]
        Measurement = "msexchange.rpcclient"
```

**MySQL**

Here are the configuration details:

```
[[inputs.mysql]]


 ## specify servers via a url matching:
 ##  [username[:password]@][protocol[(address)]]/[?tls=[true|false|skip-verify]]
 ##  see https://github.com/go-sql-driver/mysql#dsn-data-source-name
 ##  e.g.
 ##    servers = ["user:passwd@tcp(127.0.0.1:3306)/?tls=false"]
 ##    servers = ["user@tcp(127.0.0.1:3306)/?tls=false"]
 #
 ## If no servers are specified, then localhost is used as the host.
 ##servers = ["MS_USER:MS_PASSWORD@tcp(localhost:MS_PORT)/?tls=false"]
 servers = ["root:VMware1!@tcp(localhost:3306)/?tls=false"]
 ## the limits for metrics form perf_events_statements
 perf_events_statements_digest_text_limit  = 120
 perf_events_statements_limit              = 250
 perf_events_statements_time_limit         = 86400
 #
 ## if the list is empty, then metrics are gathered from all databasee tables
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
#
## gather metrics from INFORMATION_SCHEMA.TABLES for databases provided above list
gather_table_schema                     = true
#
## gather thread state counts from INFORMATION_SCHEMA.PROCESSLIST
gather_process_list                     = true
#
## gather auto_increment columns and max values from information schema
gather_info_schema_auto_inc             = true
#
## gather metrics from SHOW SLAVE STATUS command output
gather_slave_status                     = true
#

## gather metrics from SHOW BINARY LOGS command output

  gather_binary_logs                      = false

## gather thread state counts from INFORMATION_SCHEMA.USER_STATISTICS
gather_user_statistics                  = false
#
## gather metrics from PERFORMANCE_SCHEMA.TABLE_IO_WAITS_SUMMARY_BY_TABLE
gather_table_io_waits                   = true
#
## gather metrics from PERFORMANCE_SCHEMA.TABLE_LOCK_WAITS
gather_table_lock_waits                 = true
#
## gather metrics from PERFORMANCE_SCHEMA.TABLE_IO_WAITS_SUMMARY_BY_INDEX_USAGE
gather_index_io_waits                   = true
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
 #
 ## gather metrics from PERFORMANCE_SCHEMA.EVENT_WAITS
 gather_event_waits                      = true
 #
 ## gather metrics from PERFORMANCE_SCHEMA.FILE_SUMMARY_BY_EVENT_NAME
 gather_file_events_stats                = true
 #
 ## gather metrics from PERFORMANCE_SCHEMA.EVENTS_STATEMENTS_SUMMARY_BY_DIGEST
 gather_perf_events_statements           = true
 #
 ## Some queries we may want to run less often (such as SHOW GLOBAL VARIABLES)
 interval_slow                   = "30m"


 ## Optional SSL Config

   # ssl_ca = "/etc/telegraf/ca.pem"




   # ssl_cert = "/etc/telegraf/cert.pem"




   # ssl_key = "/etc/telegraf/key.pem"


 fieldpass = ["aborted_clients", "bytes_sent", "bytes_received", "binlog_cache_size",
"connection_errors_accept", "connection_errors_internal",
"connection_errors_max_connections", "queries", "threads_cached", "threads_connected",
"threads_running", "uptime", "delayed_insert_limit", "delayed_insert_timeout",
"delayed_queue_size", "max_connect_errors", "max_connections", "max_delayed_threads",
"max_error_count", "binary_size_bytes", "binary_files_count",
"commands_show_processlist", "access_denied", "bytes_received", "bytes_sent",
"commit_transactions", "concurrent_connections", "connected_time", "denied_connections",
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
"denied_connections", "rollback_transactions", "rows_fetched", "rows_updated",
"select_commands", "total_connections", "table_io_waits_total_fetch",
"table_io_waits_total_insert", "table_io_waits_total_update",
"table_io_waits_total_delete", "connections", "aborted_connects",
"events_waits_seconds_total", "events_waits_total", "innodb_print_all_deadlocks",
"innodb_open_files", "innodb_buffer_pool_size", "innodb_row_lock_time_avg",
"innodb_row_lock_current_waits", "innodb_row_lock_time_max", "innodb_table_locks",
"innodb_row_lock_waits", "innodb_buffer_pool_dump_status",
"innodb_buffer_pool_load_status", "innodb_buffer_pool_pages_data",
"innodb_buffer_pool_bytes_data", "innodb_buffer_pool_pages_dirty",
"innodb_buffer_pool_bytes_dirty", "innodb_buffer_pool_pages_flushed",
"innodb_checksums", "events_statements_no_index_used_total",
"*io_waits_seconds_total_fetch", "read_high_priority", "write_concurrent_insert"]
```

**NGINX**

Here are the configuration details:

```
# Read Nginx's basic status information (ngx_http_stub_status_module)

[[inputs.nginx]]


  ## An array of Nginx stub_status URI to gather stats.

  urls = ["http://127.0.0.1:8081/nginx_status"]


  ## Optional SSL Config

    # ssl_ca = "/etc/telegraf/ca.pem"

    # ssl_cert = "/etc/telegraf/cert.pem"

    # ssl_key = "/etc/telegraf/key.pem"


  ## Use SSL but skip chain & host verification

    # insecure_skip_verify = false


  ## HTTP response timeout (default: 5s)

  response_timeout = "5s"
```

**NTPD**

Here are the configuration details:

```
[[inputs.ntpq]]

     name_prefix = "ntpd."

     ## If false, set the -n ntpq flag. Can reduce metric gather times.
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
    dns_lookup = tr
```

**OracleDB**

- Complete the prerequisites mentioned in the Pre-Requirements for Application Services.
- Download `wavefront_oracle_metrics.py` from cloud proxy to the Telegraf configuration directory. For example, `C:\Telegraf\telegraf-1.20.4\telegraf.d`.
  - Invoke-WebRequest `https://$CP_IP/downloads/salt/content-accessories/wavefront_oracle_metrics.py -OutFile C:\Telegraf\telegraf-1.20.4\telegraf.d\wavefront_oracle_metrics.py`

Here are the configuration details:

```
# Read metrics exposed by chef

[[inputs.exec]]



  # commands = ['python "wavefront_oracle_metrics.py script path" -u "OracleDB Username"
-p "OracleDB Password" -s "OracleDB SID/DSN"']

  commands = ['python "C:
\Telegraf\telegraf-1.20.4\telegraf.d\wavefront_oracle_metrics.py" -u "orcl.vrops.es" -p
"Qw123456" -s "https://VAP-test.vrops.es:1158/em"']



  timeout = "180s"

  data_format = "influx"

  name_prefix = "oracledb."
```

**PostgreSQL**

Here are the configuration details:

```
# Read metrics from one or many postgresql servers

[[inputs.postgresql]]



  ## specify address via a url matching:

  ##   postgres://[pqgotest[:password]]@localhost[/dbname]\

  ##       ?sslmode=[disable|verify-ca|verify-full]

  ## or a simple string:

  ##   host=localhost user=pqotest password=... sslmode=... dbname=app_production

  ##

  ## All connection parameters are optional.

  ##

  ## Without the dbname parameter, the driver will default to a database
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  ## with the same name as the user. This dbname is just for instantiating a
  ## connection with the server and doesn't restrict the databases we are trying
  ## to grab metrics for.
  ##
  address = "host=localhost user=postgres password=\t Password1! port=5432"
  ##address = "host=localhost user=postgres  password=\t Password1! port=5432"
  ## A  list of databases to explicitly ignore.  If not specified, metrics for all
  ## databases are gathered.  Do NOT use with the 'databases' option.
    # ignored_databases = ["postgres", "template0", "template1"]



  ## A list of databases to pull metrics about. If not specified, metrics for all
  ## databases are gathered.  Do NOT use with the 'ignored_databases' option.

    # databases = ["app_production", "testing"]

  ## Optional SSL Config
    # ssl_ca = "/etc/telegraf/ca.pem"
    # ssl_cert = "/etc/telegraf/cert.pem"
    # ssl_key = "/etc/telegraf/key.pem"

  ## Use SSL but skip chain & host verification

    # insecure_skip_verify = false
```

**Pivotal**

Here are the configuration details:

```
# Read Tomcat status information (mod_status)
[[inputs.jolokia2_agent]]
  ## urls = ["http://localhost:8080/jolokia"]
  urls = ["http://localhost:8080/jolokia"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
name_prefix = "pivotalserver."


## Optional SSL Config
  # ssl_ca = "/etc/telegraf/ca.pem"
  # ssl_cert = "/etc/telegraf/cert.pem"
  # ssl_key = "/etc/telegraf/key.pem"


## Use SSL but skip chain & host verification
  # insecure_skip_verify = false


### JVM Generic


[[inputs.jolokia2_agent.metric]]
  name = "OperatingSystem"
  mbean = "java.lang:type=OperatingSystem"
  paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]


[[inputs.jolokia2_agent.metric]]
  name = "OperatingSystem"
  mbean = "java.lang:type=OperatingSystem"
  paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]


[[inputs.jolokia2_agent.metric]]
  name  = "jvm_runtime"
  mbean = "java.lang:type=Runtime"
  paths = ["Uptime"]


[[inputs.jolokia2_agent.metric]]
  name  = "jvm_memory"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
    mbean = "java.lang:type=Memory"
    paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]


  [[inputs.jolokia2_agent.metric]]
    name     = "jvm_garbage_collector"
    mbean    = "java.lang:name=*,type=GarbageCollector"
    paths    = ["CollectionTime", "CollectionCount"]
    tag_keys = ["name"]


  [[inputs.jolokia2_agent.metric]]
    name       = "jvm_memory_pool"
    mbean      = "java.lang:name=*,type=MemoryPool"
    paths      = ["Usage", "PeakUsage", "CollectionUsage"]
    tag_keys   = ["name"]
    tag_prefix = "pool_"


  ### TOMCAT
  [[inputs.jolokia2_agent.metric]]
    name     = "GlobalRequestProcessor"
    mbean    = "Catalina:name=*,type=GlobalRequestProcessor"
    paths    =
["requestCount","bytesReceived","bytesSent","processingTime","errorCount"]
    tag_keys = ["name"]


  [[inputs.jolokia2_agent.metric]]
    name     = "JspMonitor"
    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,name=jsp,type=JspMonitor"
    paths    = ["jspReloadCount","jspCount","jspUnloadCount"]
    tag_keys = ["J2EEApplication","J2EEServer","WebModule"]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  [[inputs.jolokia2_agent.metric]]
    name     = "ThreadPool"
    mbean    = "Catalina:name=*,type=ThreadPool"
    paths    = ["maxThreads","currentThreadCount","currentThreadsBusy"]
    tag_keys = ["name"]


  [[inputs.jolokia2_agent.metric]]
    name     = "Servlet"
    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,j2eeType=Servlet,name=*"
    paths    = ["processingTime","errorCount","requestCount"]
    tag_keys = ["name","J2EEApplication","J2EEServer","WebModule"]


  [[inputs.jolokia2_agent.metric]]
    name     = "Cache"
    mbean    = "Catalina:context=*,host=*,name=Cache,type=WebResourceRoot"
    paths    = ["hitCount","lookupCount"]
    tag_keys = ["context","host"]


  [[inputs.jolokia2_agent.metric]]
    name  = "java_class_loading"
    mbean = "java.lang:type=ClassLoading"
    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]


  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths = ["ThreadCount"]
```

*Table continued on next page*

*Continued from previous page*

| Configuration Details for Supported Application Services |
| --- |
| [[inputs.jolokia2_agent.metric]]<br>name = "java_buffer_pool"<br>mbean = "java.nio:name=*,type=BufferPool"<br>paths = ["Count", "MemoryUsed", "TotalCapacity"]<br>tag_keys = ["name"] |
| **RabbitMQ**<br>Here are the configuration details:<br># Read metrics from one or many RabbitMQ servers via the management API<br>[[inputs.rabbitmq]]<br>arc_service_id = "145d46c3-3f5d-449b-af93-f32a150c4965"<br><br>##url = "http://localhost:15672"<br>url = "http://10.218.135.219:15672"<br><br>metric_exclude = ["federation"]<br><br># name = "rmq-server-1" # optional tag<br>username = "admin"<br>password = "admin"<br><br>## Optional SSL Config<br># ssl_ca = "/etc/telegraf/ca.pem"<br># ssl_cert = "/etc/telegraf/cert.pem"<br># ssl_key = "/etc/telegraf/key.pem"<br><br>## Use SSL but skip chain & host verification<br># insecure_skip_verify = false<br><br>## Optional request timeouts |

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  ##
  ## ResponseHeaderTimeout, if non-zero, specifies the amount of time to wait
  ## for a server's response headers after fully writing the request.
  # header_timeout = "3s"
  ##
  ## client_timeout specifies a time limit for requests made by this client.
  ## Includes connection time, any redirects, and reading the response body.
  # client_timeout = "4s"


  ## A list of nodes to pull metrics about. If not specified, metrics for
  ## all nodes are gathered.
  # nodes = ["rabbit@node1", "rabbit@node2"]
```

**Riak**

Here are the configuration details:

```
# Read Riak status information (mod_status)
[[inputs.riak]]


  # Specify a list of one or more riak http servers
  ##servers = ["http://localhost:8098"]


  servers = ["http://127.0.0.1:8098"]
```

**SharePoint Server**

Here are the configuration details:

```
[[inputs.win_services]]
  name_prefix = "sharepoint."


   service_names = [
   "OSearch15",
   "SPAdminV4",
   "SPSearchHostController",
   "SPTimerV4",
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
   "SPTraceV4",
   "SPUserCodeV4",
   "SPWriterV4"
]


[[inputs.win_perf_counters]]
  plugin_name_override="sharepoint"


  [[inputs.win_perf_counters.object]]
    ObjectName = "Process"
    Instances = ["mssearch", "WSSADMIN", "hostcontrollerservice", "OWSTIMER",
"wsstracing", "SPUCHostService", "SPWRITER"]
    Counters = ["% Processor Time", "Thread Count"]
    Measurement = "sharepoint.process"


  [[inputs.win_perf_counters.object]]
    ObjectName = "Web Service"
    Instances = ["*"]
    Counters = ["Bytes Sent/sec", "Bytes Received/sec"]
    Measurement = "sharepoint.webservice"


  [[inputs.win_perf_counters.object]]
    ObjectName = "SharePoint Foundation"
    Counters = ["Object Cache Hit %", "Executing Sql Queries", "Executing Time/Page
Request", "Reject Page Requests Rate", "Incoming Page Requests Rate", "Active Threads",
"Object Cache Hit Count", "Current Page Requests", "Sql Query Executing  time",
"Responded Page Requests Rate"]
    Instances = ["_total"]
    Measurement = "sharepoint.Foundation"


  [[inputs.win_perf_counters.object]]
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
    ObjectName = "SharePoint Records Management Counters"

    Counters = ["Search results processed / sec"]

    Instances = ["------"]

    Measurement = "sharepoint.RecordsManagementCounters"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation Search Gatherer Projects"

    Counters = ["Crawls in progress", "Filtered Text Rate", "Filtered Office Rate",
"Filtered HTML Rate", "Accessed File Rate", "Accessed HTTP Rate", "File Errors Rate",
"HTTP Errors Rate"]

    Instances = ["_total"]

    Measurement = "sharepoint.FoundationSearchGathererProjects"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation Search Schema Plugin"

    Counters = ["Total Documents"]

    Instances = ["_total"]

    Measurement = "sharepoint.FoundationSearchSchemaPlugin"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation BDC Online"

    Counters = ["Total calls failed", "Total calls per second"]

    Instances = ["------"]

    Measurement = "sharepoint.FoundationBDCOnline"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation Search Gatherer"

    Counters = ["Changes Processed", "Threads Committing Transactions", "Time Outs",
"Active Queue Length", "Idle Threads", "Heartbeats Rate"]

    Instances = ["------"]

    Measurement = "sharepoint.FoundationSearchGatherer"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation Search Indexer Plugin"

    Counters = ["Propagation Rate", "Average Query Latency", "Queries Succeeded",
"Queries Failed", "Queries"]

    Instances = ["_total"]

    Measurement = "sharepoint.FoundationSearchIndexerPlugin"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation Search Query Processor"

    Counters = ["Security Descriptor Cache Misses"]

    Instances = ["------"]

    Measurement = "sharepoint.FoundationSearchQueryProcessor"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation Search FAST Content Plugin"

    Counters = ["Batches Failed Timeout", "Submission Timeouts", "Items Failed Total",
"Items Failed Timeout"]

    Instances = ["------"]

    Measurement = "sharepoint.FoundationSearchFASTContentPlugin"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation Search Archival Plugin"

    Counters = ["Queues Committing", "Queues Waiting", "Queues Filtering", "Queues
Available"]

    Instances = ["_total"]

    Measurement = "sharepoint.FoundationSearchArchivalPlugin"


  [[inputs.win_perf_counters.object]]

    ObjectName = "SharePoint Foundation BDC Metadata"

    Counters = ["Cache misses per second", "Cache hits per second"]
```

*Table continued on next page*

*Continued from previous page*

| **Configuration Details for Supported Application Services** |
| --- |
| ```<br>    Instances = ["_total"]<br>    Measurement = "sharepoint.FoundationBDCMetadata"<br><br>  [[inputs.win_perf_counters.object]]<br>    ObjectName = "SharePoint Foundation Search Gatherer Databases"<br>    Counters = ["Documents in the crawl history", "Documents in the crawl queue"]<br>    Instances = ["_total"]<br>    Measurement = "sharepoint.FoundationSearchGathererDatabases"<br>``` |
| **Tomcat Server**<br>Here are the configuration details:<br>```<br># Read Tomcat status information (mod_status)<br>[[inputs.jolokia2_agent]]<br>  ##urls = ["http://localhost:8080/jolokia"]<br>  urls = ["http://localhost:8083/jolokia"]<br>  name_prefix = "tomcat."<br><br>  ## Optional SSL Config<br>    # ssl_ca = "/etc/telegraf/ca.pem"<br>    # ssl_cert = "/etc/telegraf/cert.pem"<br>    # ssl_key = "/etc/telegraf/key.pem"<br><br>  ## Use SSL but skip chain & host verification<br>    # insecure_skip_verify = false<br><br>  ### JVM Generic<br><br>  [[inputs.jolokia2_agent.metric]]<br>    name = "OperatingSystem"<br>    mbean = "java.lang:type=OperatingSystem"<br>    paths = ["ProcessCpuLoad","SystemLoadAverage","SystemCpuLoad"]<br>``` |

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
[[inputs.jolokia2_agent.metric]]
  name = "OperatingSystem"
  mbean = "java.lang:type=OperatingSystem"
  paths = ["MaxFileDescriptorCount","OpenFileDescriptorCount"]

[[inputs.jolokia2_agent.metric]]
  name  = "jvm_runtime"
  mbean = "java.lang:type=Runtime"
  paths = ["Uptime"]

[[inputs.jolokia2_agent.metric]]
  name  = "jvm_memory"
  mbean = "java.lang:type=Memory"
  paths = ["HeapMemoryUsage", "NonHeapMemoryUsage", "ObjectPendingFinalizationCount"]

[[inputs.jolokia2_agent.metric]]
  name     = "jvm_garbage_collector"
  mbean    = "java.lang:name=*,type=GarbageCollector"
  paths    = ["CollectionTime", "CollectionCount"]
  tag_keys = ["name"]

[[inputs.jolokia2_agent.metric]]
  name       = "jvm_memory_pool"
  mbean      = "java.lang:name=*,type=MemoryPool"
  paths      = ["Usage", "PeakUsage", "CollectionUsage"]
  tag_keys   = ["name"]
  tag_prefix = "pool_"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
  ### TOMCAT

  [[inputs.jolokia2_agent.metric]]

    name     = "GlobalRequestProcessor"

    mbean    = "Catalina:name=*,type=GlobalRequestProcessor"

    paths    =
["requestCount","bytesReceived","bytesSent","processingTime","errorCount"]

    tag_keys = ["name"]


  [[inputs.jolokia2_agent.metric]]

    name     = "JspMonitor"

    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,name=jsp,type=JspMonitor"

    paths    = ["jspReloadCount","jspCount","jspUnloadCount"]

    tag_keys = ["J2EEApplication","J2EEServer","WebModule"]


  [[inputs.jolokia2_agent.metric]]

    name     = "ThreadPool"

    mbean    = "Catalina:name=*,type=ThreadPool"

    paths    = ["maxThreads","currentThreadCount","currentThreadsBusy"]

    tag_keys = ["name"]


  [[inputs.jolokia2_agent.metric]]

    name     = "Servlet"

    mbean    =
"Catalina:J2EEApplication=*,J2EEServer=*,WebModule=*,j2eeType=Servlet,name=*"

    paths    = ["processingTime","errorCount","requestCount"]

    tag_keys = ["name","J2EEApplication","J2EEServer","WebModule"]


  [[inputs.jolokia2_agent.metric]]

    name     = "Cache"

    mbean    = "Catalina:context=*,host=*,name=Cache,type=WebResourceRoot"
```

*Table continued on next page*

*Continued from previous page*

**Configuration Details for Supported Application Services**

```
    paths    = ["hitCount","lookupCount"]
    tag_keys = ["context","host"]


  [[inputs.jolokia2_agent.metric]]
    name  = "java_class_loading"
    mbean = "java.lang:type=ClassLoading"
    paths = ["LoadedClassCount", "UnloadedClassCount", "TotalLoadedClassCount"]


  [[inputs.jolokia2_agent.metric]]
    name  = "java_threading"
    mbean = "java.lang:type=Threading"
    paths = ["ThreadCount"]


  [[inputs.jolokia2_agent.metric]]
    name  = "java_buffer_pool"
    mbean = "java.nio:name=*,type=BufferPool"
    paths = ["Count", "MemoryUsed", "TotalCapacity"]
    tag_keys = ["name"]
```

Oracle WebLogic

**WebSphere**

Here are the configuration details:

```
[[inputs.httpjson]]
  servers = ["http://localhost:9080/wasmonitor/stats"]
  method = "POST"
  name_override="websphere"


  response_timeout = "60s"
  tag_keys = ["name", "process", "node", "cell", "mbeanIdentifier"]
  [inputs.httpjson.headers]
    Authorization = "Basic d2F2ZWZyb250OndhdmVmcm9udA=="
```

*Table continued on next page*

*Continued from previous page*

| Configuration Details for Supported Application Services |
| --- |
| [[inputs.exec]]<br>commands = ['C:\\VMware\\UCP\\ucp-minion\\bin\\run-py-script.bat "C:\VMware\UCP\ucp-minion\lib\site-packages\ucp\application_availability.py" --config="C:\VMware\UCP\salt\conf\grains" "java" "http://localhost:9080"']<br>timeout = "180s"<br>data_format = "influx"<br>name_prefix = "websphere."<br>[inputs.exec.tags]<br>server = "http://localhost:9080/wasmonitor/stats" |

**Table 93: Configuration Details for Custom Monitoring Plugins**

| |
| --- |
| **Ping check**<br>Here are the configuration details:<br>[[inputs.ping]]<br>## Hosts to send ping packets to.<br>urls = [www.vmware.com]<br><br>## Number of ping packets to send per interval.  Corresponds to the "-c"<br>## option of the ping command.<br># count = 1<br><br>## Time to wait between sending ping packets in seconds.  Operates like the<br>## "-i" option of the ping command.<br># ping_interval = 1.0<br><br>## If set, the time to wait for a ping response in seconds.  Operates like<br>## the "-W" option of the ping command.<br># timeout = 1.0 |

*Table continued on next page*

*Continued from previous page*

```
## If set, the total ping deadline, in seconds.  Operates like the -w option
## of the ping command.
# deadline = 10


## Interface or source address to send ping from.  Operates like the -I or -S
## option of the ping command.
# interface = ""
```

**TCP check**

Here are the configuration details:

```
[[inputs.net_response]]
name_override="netresponse"
## Server address (default localhost)
address = "10.10.10.10:443"
protocol = "tcp"


## Set timeout
# timeout = "1s"


## Set read timeout (only used if expecting a response)
# read_timeout = "1s"
```

**HTTP check**

Here are the configuration details:

```
# HTTP/HTTPS request given an address a method and a timeout
[[inputs.http_response]]
name_override="httpresponse"
## List of url to query.
address = https://10.10.10.10/suite-api/api/adapters
method = "GET"


## Set http_proxy.
## Telegraf uses the system wide proxy settings if it's is not set.
```

*Table continued on next page*

*Continued from previous page*

```
# http_proxy = http://localhost:8888


## Set response_timeout (default 5 seconds)

# response_timeout = "5s"


## Whether to follow redirects from the server (defaults to false)

# follow_redirects = false


## Optional name of the field that will contain the body of the response.

## By default it is set to an empty String indicating that the body's content won't be
added

# response_body_field = ''


## Optional substring or regex match in body of the response (case sensitive)

# response_string_match = "\"service_status\": \"up\""

# response_string_match = "ok"

# response_string_match = "\".*_status\".?:.?\"up\""




## Optional SSL Config

# ssl_ca = "/etc/telegraf/ca.pem"

# ssl_cert = "/etc/telegraf/cert.pem"

# ssl_key = "/etc/telegraf/key.pem"




## Use SSL but skip chain & host verification

insecure_skip_verify = true


## HTTP Request Headers (all values must be strings)

[inputs.http_response.headers]
```

*Table continued on next page*

*Continued from previous page*

```
accept = "application/json" Authorization = "OpsToken a16f7a2b-
b033-48bc-9827-2daf8e205537::ec11ee5f-8623-4558-a904-8b4ea3f6f47d"
```

**Custom scripts**

Here are the configuration details:

**For Linux Platforms:**

```
# Read metrics from command that can output to stdout

[[inputs.exec]]

name_prefix = "executescript."

name_override = "output"

# commands = ["<prefix> <file path> <argument>"]

commands = ["python /opt/scripts/argument.py 10"]


## Data format to consume.

data_format = "value"

data_type = "integer"


## Timeout for each command to complete.

timeout = "300s"


[inputs.exec.tags]


file_path = "/opt/scripts/argument.py"

script_name = "my-py-script on Centos7-VM"
```

**For Windows Platforms:**

```
# Read metrics from command that can output to stdout

[[inputs.exec]]


name_prefix = "executescript."

name_override = "output"


# commands = ["<prefix> <file path> <argument>"]
```

*Table continued on next page*

*Continued from previous page*

```
commands = ["Powershell -File C:\\\\VMware\\\\Scripts\\\\Arguments.ps1 vmtoolsd"]


## Data format to consume.
data_format = "value"
data_type = "integer"


## Timeout for each command to complete.
timeout = "300s"


[inputs.exec.tags]


file_path = "C:\\VMware\\Scripts\\Arguments.ps1"
script_name = "my-script-for-vmtool on Windows2022-PhysicalServer"
```

**Linux Process**

Here are the configuration details (via regular expression):

```
# Monitor process cpu and memory usage
[[inputs.procstat]]
  ## pattern as argument for pgrep (ie, pgrep -f <pattern>)
  pattern = ".*ucp-.*"


  fieldpass =  ["running", "cpu_usage", "memory_usage"]
  [inputs.procstat.tags]
    search_pattern = "regex_.#!AsTeRiSk!#ucp-.#!AsTeRiSk!#"
```

Here are the configuration details (via executable name):

```
# Monitor process cpu and memory usage
[[inputs.procstat]]
  ## executable name (ie, pgrep <exe>)
  exe = "top"
  fieldpass =  ["running", "cpu_usage", "memory_usage"]
  [inputs.procstat.tags]
```

*Table continued on next page*

*Continued from previous page*

```
    search_pattern = "exec_top"
```

Here are the configuration details (via process ID):

```
# Monitor process cpu and memory usage
[[inputs.procstat]]
  ## PID file to monitor process
  pid_file = "/var/run/vmtoolsd.pid"

  fieldpass =  ["running", "cpu_usage", "memory_usage"]
  [inputs.procstat.tags]

    search_pattern = "pidfile_/var/run/vmtoolsd.pid"
```

**Windows Service**

Here are the configuration details (via Windows service name):

```
# Monitor process cpu and memory usage
[[inputs.procstat]]

## Windows service name
win_service = "Dhcp"
pid_finder = "native"

fieldpass =  ["running", "cpu_usage", "memory_usage"]
[inputs.procstat.tags]
search_pattern = "exec_Dhcp"

[[inputs.win_services]]
service_names = ["Dhcp"]
name_override="procstat"

[inputs.win_services.tags]
```

*Table continued on next page*

*Continued from previous page*

```
search_pattern = "exec_Dhcp"
```

**Troubleshooting (Open Source Telegraf)**

**Windows Services**

If you configure multiple services and the service names include underscores, then the associated objects have the same display name.

While monitoring a Windows Services plugin using open source Telegraf, if there are multiple services with names such as, `svc1_name1` and `svc1_name2` with an underscore in the service names, then both the objects have the same display name.

If you want the display names to have different names, follow these steps.

1. Manually edit the service name in VMware Cloud Foundation OperationsVMware Cloud Foundation Operations after the Windows Services object is created so that the correct service name is displayed as the display name of the Service object.
2. Navigate to **Operations › Configurations**, and then click the **Inventory Management** tile. From the **Objects** tab, select the Service object with the incorrect name and edit it.
3. Add the value as entered in the **Service Name** field to the **Display Name** field.

   For example, while monitoring a Windows service with the service name `SMS_SITE_COMPONENT_MANAGER`, if the display name of the current object is `Services exec SMS on Windows_OS_on_<windows hostname>`, where `SMS` is the display name, change the display name to `Services exec SMS_SITE_COMPONENT_MANAGER on Windows_OS_on_<windows hostname>`.

4. Click **OK**.

## Monitoring Physical Servers

If you want to start monitoring physical servers, use product-managed Telegraf agents. Use the helper script to install the product-managed agents. For more information, see Install/Uninstall an Agent Using a Script on a Linux Platform and Install/Uninstall an Agent Using a Script on a Windows PlatformInstall/Uninstall an Agent Using a Script on Linux Platforms and Install/Uninstall an Agent Using a Script on a Windows Platform . Existing configurations for user-managed Telegraf agents that are used to monitor physical servers is supported, data collection will continue, and additional steps are not required.

## Service and Application Discovery

You can discover services and applications using the Service Discovery adapter.

**Service Discovery**

Service discovery helps you discover services running in each VM and then builds a relationship or dependency between the services from different VMs. You can view basic metrics based on the services you want to monitor. You can also use the service discovery dashboards to monitor the services.

Service discovery helps you determine the kind of services running on each VM in your environment. You can find out which VM is a part of a service, the impact of shutting down or moving a VM, the impact of an incident, and the right

escalation path for a problem. You can also determine which VMs are used to migrate a service and which services are impacted by a planned outage on a VM or an infrastructure component.

**Application Discovery**

Applications can also be discovered on a VM with services connected to each other and those that talk to each other. You can discover predefined and custom applications. VMware Cloud Foundation Operations for logs and VMware Aria Operations are the predefined applications. You can view the VMs connected to their services, how the services talk to each other, and those that are connected to different VMs.

### Licensing

You can discover and monitor services using VMware Aria Operations Advanced and Enterprise editions.

To discover and monitor services and add and view applications, follow these steps in VMware Aria OperationsVMware Cloud Foundation Operations:

- Configure Service and Application Discovery. For more information, see Configure Service and Application Discovery.
- Manage Services and View Applications. For more information, see Manage Services and View Applications.
- Monitor services using dashboards. For more information, see Service Discovery Dashboards.
- View the services discovered. For more information, see Discovered Services.

### Supported Platforms and Products for Service Discovery

Service discovery supports specific platforms and product versions.

You can either provide guest operating system credentials with appropriate privileges or use the credential-less approach to discover services.

**Supported Product Versions for Credential-Based Service Discovery**

- For ESXi, vCenter, and VMware Cloud on AWS versions, see the VMware Product Interoperability Matrix.
- VMware Tools: For details, see KB 75122.

**Supported Product Versions and Other Pre-Requisites for Credential-Less Service Discovery**

For information, see KB 78216.

**Operating System Versions**

| Operating Systems | Version |
|---|---|
| Windows | Windows 7, Windows Server 2008/R2, and above. |
| Linux | Photon, RHEL, CentOS, SUSE Linux Enterprise Server, OEL, and Ubuntu (all Linux operating systems must be based on kernel version 2.6.25 or above). |

### Supported Services

Service and application discovery supports several services that are supported in VMware Aria OperationsVMware Cloud Foundation Operations. The supported services are listed here.

**Supported Services:**

- Active Directory
- Apache HTTP

- Apache Tomcat
- Cassandra
- DB2
- Exchange Client Access Server
- Exchange Edge Transport Server
- Exchange Hub Transport Server
- Exchange Mailbox Server
- Exchange Server
- Exchange Unified Messaging Server
- GemFire
- IIS
- JBoss
- MS-SQL DB
- MySQL DB
- Nginx
- Oracle DB
- RabbitMQ
- SharePoint
- SharePoint Application Server
- SharePoint Server
- SharePoint Web Server
- SRM vCenter Replication Management Server
- SRM vCenter Replication Server
- Sybase DB
- tc Server
- vCenter Site Recovery Manager Server
- vCloud Director
- VMware vCenter
- VMware vCenter (Appliance)
- VMware View Server
- vROps Analytics
- vROps Collector
- vROps GemFire
- vROps Postgres Data
- vROps Postgres Repl
- vROps UI
- vRLI Daemon
- vRLI vInternalization
- vRLI UI
- WebLogic
- WebSphere

**Configure Service and Application Discovery**

To discover applications and services and their relationships and to access basic monitoring, you can either provide guest operating system credentials with appropriate privileges or use the credential-less approach to discover services.

- You must have a vCenter Adapter instance configured and monitoring the same vCenter that is used to discover services.
  For credential-based service discovery, the configured vCenter user must have the following privileges:

  - key: VirtualMachine.GuestOperations.ModifyAliases, Localization: Guest operations -> Guest operation alias modification
  - key: VirtualMachine.GuestOperations.QueryAliases, Localization: Guest operations -> Guest operation alias query
  - key: VirtualMachine.GuestOperations.Modify, Localization: Guest operations -> Guest operation modifications
  - key: VirtualMachine.GuestOperations.Execute, Localization: Guest operations -> Guest operation program execution
  - key: VirtualMachine.GuestOperations.Query, Localization: Guest operations -> Guest operation queries

  For credential-less service discovery, the configured vCenter user must have the following privileges:
  - key: VirtualMachine.Namespace.Management, Localization: Service Configuration -> Manage service configurations
  - key: VirtualMachine.Namespace.ModifyContent, Localization: Service Configuration -> Modify service configuration
  - key: VirtualMachine.Namespace.Query, Localization: Service Configuration -> Query service configurations
  - key: VirtualMachine.Namespace.ReadContent, Localization: Service Configuration -> Read service configuration
- The ESXi instance that hosts the VMs where services should be discovered, must have HTTPS access to port 443 from the collector node on which the service discovery adapter instance is configured.
- The ESXi instance that hosts the VMs where services should be discovered, must have HTTPS access to port 443 from the cloud proxy on which the service discovery adapter instance is configured.
- Verify that the following types of commands and utilities are used:

| Type | Commands and Utilities |
|---|---|
| UNIX Operating Systems | |
| Service Discovery | `ps`, `ss`, and `top` |
| Performance Metrics Collection | : `awk`, `csh`, `ps`, `pgrep`, and `procfs` (file system) |
| Windows Operating Systems | |
| Service Discovery | `wmic`, `netstat`, `findstr`, `net`, `reg`, and `sort`<br><br>**NOTE**<br>PowerShell Scripts are now used instead of batch scripts in case of credential-less discovery starting with VMware Tools version 12.3.0 or above. See Deprecated Features for Windows Client for more details.<br><br>**NOTE**<br>If you are using credential-less discovery and VMware Tools version 12.3.0 and above on a Windows VM, ensure that the following modules are installed on the OS::<br>• Microsoft.PowerShell.Core (default)<br>• Microsoft.PowerShell.Utility<br>• Microsoft.PowerShell.Management<br>• SmbShare (for net share)<br>• netstat utility<br>Also, ensure that the version of the OS is Windows Server 2008R2 SP2 and above. |
| Performance Metrics Collection | `wmic`, `typeperf`, and `tasklist` |

- User Access Restrictions
  - For Linux operating systems, ensure that the user is a root or member of the *`sudo`* users group.

**NOTE**
For non-root users, the
`NOPASSWD`
option must be activated in `/etc/sudoers` file to avoid the metrics collector scripts from waiting for the interactive password input.
Steps to activate the
`NOPASSWD`
option for a particular sudo user:

1. Login to the specific VM as a root user.
2. Run the **`sudo visudo`** command that opens an editor.
3. In the command section, add username ALL=(ALL) NOPASSWD:<ss path>, <awk path>, <netstat path>. The username must be replaced with an existing user name for which this option is activated. Example: vmware ALL=(ALL) NOPASSWD: `/usr/sbin/ss`, `/usr/bin/netstat`, `/user/bin/awk`.
   When you perform the Execute Script action and you need to use `command/utilities`, for those commands that need a sudo user password provision, the full path of `command/utility` must be added to the NOPASSWD commands list.
4. Save the file and close it. It is automatically reloaded.

– To discover services on Windows, the local administrator account must be configured.

**NOTE**
Services will not be discovered for administrator group members that are different from the administrator account itself if the policy setting
`User Account Control: Run all administrators in Admin Approval Mode`
is turned on. As a workaround, you can turn off this policy setting to discover services. However, if you turn the policy setting off, the security of the operating system is reduced.

– To discover services on Windows Active Directory, the domain administrator account must be configured.

- The system clock must be synchronized between the VMware Aria OperationsVMware Cloud Foundation Operations nodes, the vCenter, and the VM if service discovery is working in credential-based mode and guest alias mapping is used for authentication.
- The configured user must have read and write privileges to the temp directory (execute privilege is also required on this directory in Linux systems). For Windows systems, the path can be taken from the environment variable *`TEMP`*. For Linux systems, it is `/tmp` and/or `/var/tmp`.
- The SSO Server URL must be reachable from the VMware Aria OperationsVMware Cloud Foundation Operations node on which the service discovery adapter is located.
- For more information about supported platforms and versions, see Supported Platforms and Products for Service Discovery.

**NOTE**
If more than one VMware Aria OperationsVMware Cloud Foundation Operations instance is monitoring the same vCenter and service discovery is activated for those VMware Aria OperationsVMware Cloud Foundation Operations instances, then service discovery might be unstable, which is a known VMware Tools problem. As a result, guest operations might fail to execute.

1. From the left menu, click **Operations › Configurations** , and then from the right panel, click the **Service Discovery** tile.
2. From the **Service Discovery** page, click the **Configure Service Discovery** option.
3. From the **Integrations** page, click the vCenter instance from the list and then select the **Service Discovery** tab.

4. To activate service discovery in this vCenter, activate the **Service Discovery** option.
5. To activate application discovery in this vCenter, select the **Enable Application Discovery** check box.
6. You can choose to add credentials by selecting the **Use alternate credentials** check box.
   a) Click the plus sign and enter the details in the **Manage Credentials** dialog box, which include a credential name and a vCenter user name and password. In addition, enter the user name and password for Windows, Linux, and SRM and click **OK**.
7. Alternatively, if you are using the default user name and password, enter a default user name and password for Windows, Linux, and SRM.
8. Enter a password for the guest user mapping.
9. You can also activate grouping of the application, creation of a business application, and activate application discovery.
10. Click **Save**.

    **NOTE**
    If you specify a non-root user for Linux, services are not discovered unless you activate the option Use Sudo (Linux Non-root user) while editing the associated Service Discovery adapter instance after you create the vCenter Cloud Account. This option is deactivated by default, which means the root user is expected by default when you configure the vCenter Cloud Account.

11. Edit the cloud account created for service discovery.
12. In the **Advanced Settings** section, activate the **Application Discovery** field to discover predefined and custom applications.
13. In the **Advanced Settings** section, to configure credential-less service discovery, select **Activated** from the **Credential-less service discovery status** field.
14. Click **Save**.

You can manage services supported by VMware Aria OperationsVMware Cloud Foundation Operations on specific VMs.

### Manage Services

You can manage services supported by VMware Aria OperationsVMware Cloud Foundation Operations on the specific VMs.

**Where You Manage Services**

From the left menu, select **Operations › Applications**. From the **Applications** panel, select **Manage SDMP Services**. You can also navigate to the
`Manage SDMP Services`
tab by selecting **Operations › Configurations**, and then from the right panel, select the **Service Discovery** tile. Select the
`Manage Services`
option from the **Service Discovery** page.

You can view specific details from the options in the data grid.

**Table 94: Datagrid Options**

| Options | Description |
|---|---|
| VM Name | Name of the VM. |
| Operating System | Operating system installed on the VM. |

*Table continued on next page*

*Continued from previous page*

| Options | Description |
|---|---|
| Services Discovered | Displays the names of discovered services or `None`, if services are not discovered on the VM. |
| Service Monitoring | Displays the current value of the VM's service monitoring setting. If set, services are discovered and service performance metrics are calculated every 5 minutes. Otherwise, service discovery is performed every 24 hours. |
| Authentications Status | VM authentication status for service discovery. The possible values are:<br>• Unknown<br>• Failed<br>• Guest Alias<br>• Common Credentials<br>• Credential-less |
| Power State | Power status of the VMs. The possible values are:<br>• Powered On<br>• Powered Off<br>• Suspended<br>• Unknown |
| Collection State | Displays the collection state of an adapter instance of each object. You can see the name of the adapter instance and its state in a tool tip when you point to the collection state icon. To manage an adapter instance to start and stop collection of data, from the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile. |
| Collection Status | Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the collection status icon. To manage an adapter instance to start and stop collection of data, from the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile.<br>You can view a message for VMs with a failed authentication status in a tool tip when you point to the collection status icon. |
| vCenter Name | Name of the vCenter Adapter instance to which that VM resource belongs. |

**Table 95: Toolbar and Row Options**

| Options | Description |
|---|---|
| VM Actions | Displays a list of actions. For more information, see Actions in VMware Aria Operations. |
| **Horizontal Ellipsis › Provide Password** | Select VMs from the list, click the horizontal ellipsis and then click **Provide Password** to provide a user name and password for the selected VMs to discover the services. |

*Table continued on next page*

*Continued from previous page*

| Options | Description |
|---|---|
| **Horizontal Ellipsis › Activate Service Monitoring** | Select VMs from the list, click the horizontal ellipsis and then click **Activate Service Monitoring** to activate frequent service discovery and service performance metrics calculation (every 5 minutes).<br>**NOTE**<br>Selecting too many VMs will potentially result in vCenter degradation which is a known issue. |
| **Horizontal Ellipsis › Deactivate Service Monitoring** | Select VMs from the list, click the horizontal ellipsis and then click **Deactivate Service Monitoring** to deactivate frequent service discovery and service performance metrics calculation. Service discovery defaults to the 24-hour cycle. |
| Check box to Select/Deselect all | Selects/clears all VM object selections. |
| **Vertical Ellipsis › Go To Details** | Navigates to the **Summary** tab for the selected VM. |
| Filter | You can search through the list of VMs according to the following criteria: VM Name, Operating System, Power State, Authentication Status, and Service. |

**Create Application Definition**

Create application definitions, when you want to discover applications based on the user defined rules such as object names, tags, or properties.

1. From the left menu, click **Operations › Configurations** and then select the **Application Discovery: Rule Based** tile.
2. Click **Add Application Definition**.
3. In the **Create Application Definition** page, enter the required values.

**Table 96: Create Application Definition**

| Option | Description |
|---|---|
| Application Name | Enter a name for the application. |
| Application Prefix | Enter a prefix for the application. Prefix can help you identify the discovered applications by name. It will be prepended to the application name. |
| Select criteria for grouping discovered applications - This section allows you to define the criteria based on which the objects are grouped into an application. | |
| Group For Object Type | Select an object type to group objects by |
| Select | Select the criteria to group the objects by. The options available are:<br>• Properties<br>• Object Name<br>• Tag Category |
| Add | Click add to add another criteria |
| Reset | Click reset to clear the criteria |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Advanced Settings - This section allows you to specify information for additional filters. | |
| Select the Object Type that matches all of the following criteria | Use the drop-down list to select the object type which matches the defined criteria |
| Select | Select the object type. The options available are:<br>• Metrics<br>• Relationships<br>• Properties<br>• Object Name<br>• Tag<br>Based on your selection, you must specify the additional attributes. |
| Add | Click Add to add another criteria |
| Reset | Click Reset to clear the defined criteria |
| Add New Criteria | Click Add New Criteria to add another criteria |
| Remove Criteria | Click Remove Criteria to delete the criteria |

4. Click **Create**.
   The created application definition is displayed in the **Application Discovery: Rule Based** page.

### Manage Applications

You can discover applications based on user defined criteria such as virtual machine names, tags, or properties (rule-based applications). For example, if the application definition is assigned to a cluster compute resource called 'xyz', then the applications would be discovered only within that cluster. The object scope assignments can be made to both service based application definitions and rule based application definitions.

#### Where You Manage Applications

From the left menu, select **Operations › Applications**. From the **Applications** panel, select **Manage Applications**.

You can view direct or inherited assignments by selecting either **All** or **Self** respectively, in the sliding option at the top of the UI. You must select any of the environment objects in order to see application definitions.

You can view specific details from the options in the data grid.

**Table 97: Datagrid Options**

| Options | Description |
|---|---|
| Object Name | Name of the Object. |
| Direct Assignments | Displays the number of application definitions assigned directly to this object. |
| Inherited Assignments | Displays the number of application definitions that are inherited from the object’s ancestors. |
| Object Type | Displays the object type. |

**Table 98: Assign Application Definition To Scope of Objects**

| Options | Description |
|---|---|
| **Vertical Ellipsis** | Displays Assign Applications Definitions page. |
| Select rules from the list below to activate application discovery for selected objects | Displays the list of Application Definition Name and Application Definition Type. Select the required definition, the options available are Service-based application definition and Rule-based application definition.<br>**NOTE**<br>You can use the search option to look for application definitions name. |
| > | Expand the arrow next to the object to view the definitions assigned to the selected object and the discovered applications count. You can also view the inherited assignments for the selected object. |
| Save | Navigates to the Summary tab for the selected Object. |
| **Vertical Ellipsis › Go To Details** | Navigates to the **Summary** tab for the selected objects. |

### View Applications

You can view applications created by VMware Aria OperationsVMware Cloud Foundation Operations on specific VMs.

You can view all resources of the Application type, including:
- VMs connected to their services, how the services talk to each other and are connected to different VMs, using the Service Discovery adapter.
- Discovered applications using the VMware Aria Operations Application Management Pack.

### Where You View the Applications

From the left menu, select **Operations › Applications** to view the **Applications Home** page.

You can view the list of applications from the **Applications Home** page. Click an application to view the application in the right pane. Select an application row and then click the **Vertical Ellipsis › Go To Details** to view the object details. You can use the Filter option to search through the list of applications according to the following criteria :
- Name
- Object Type

### Options to View the Applications

After you click on an application from the **Applications Home** page, you can view details of the application or service in the right pane.

| Options | Description |
|---|---|
| View Sphere | Displays the VMs connected to their services in a spherical view. |
| View Graph | Displays the VMs connected to their services as a graph. |
| View List | Displays the VMs connected to their services in a list view. |
| View Links | You can view the links between various services. |

### Discovered Services

You can discover services using the Service Discovery adapter.

**Discovered Services**

You can view discovered services, the number of VMs on which each discovered service is running, and you can configure service discovery.

#### Where You View the Discovered Services

From the left menu, click **Operations › Configurations**. From the right panel, click the **Service Discovery** tile.

#### Discovered Services

You see a list of services that are discovered and the number of VMs that have the services running. You see this section after you have configured Service Discovery and the services are discovered.

#### Known Services

You see a list of all the services supported and those that can be discovered.

#### Custom Services

You can add a service by clicking **Add Service Definition**. You can add either a process name or Regex from the **Add Custom Service** dialog box.

**Table 99: Add Custom Service**

| Options | Description |
|---|---|
| Type | Specify the type as either process or regex.<br>**Process**:<br><br>The process name must exactly match the name that you see in the guest OS when running commands **`ps`** in Linux and **`wmic`** in Windows. Specify a single port for each service.<br><br>The following characters are not supported: `,`, `\`, and `#`.<br><br>**Regex**<br><br>Enter a regular expression that corresponds to the command line (or at least name) of the service, that you see in the guest OS when you run the following commands: `ps` in Linux and `wmic` in Windows.<br><br>For example, to discover Cassandra services, enter `cass.*dra` as the regex.<br><br>The following characters are not supported: `,` and `\n`. |

*Table continued on next page*

*Continued from previous page*

| Options | Description |
|---|---|
| | **NOTE**<br>If you specify a service using regex and the service reboots at a later time and all the ports are modified, the current service displays as offline and the new service is discovered again. If the new service has any port overlap with the previous one, a new resource will not be discovered and the previous one will continue collecting data as before. |
| Process Name | Enter a process name. |
| Port | Enter the port information. |
| Display Name | Enter the display name. |

**NOTE**
Custom service can be discovered via Service Discovery if there is a permanent listening TCP port or if there is an established UDP connection.

**Discovered Applications**

You can discover applications based on grouping the services discovered, using the Service Discovery adapter.

**Discovered Applications**
You can discover predefined applications and custom applications. VMware Cloud Foundation Operations for logs and VMware Aria Operations are the predefined applications.

**Where You View the Discovered Applications**
From the left menu, click **Operations › Configurations**, and then from the right panel click the **Application Discovery: Service Based** tile.

**Custom Applications**

You can define custom applications. Click **Add Application Definition** to add a custom application.

**Table 100: Add Custom Application**

| Option | Description |
|---|---|
| Name | Enter a name for the application. |
| Prefix | Enter a prefix for the application. |
| Services | Select a service from the list.<br>If the Service Discovery adapter discovers the service and if the services are connected to each other, a new application is discovered. The new application appears in the **Applications Home** page. Navigate to **Operations › Applications**. |

**Known Applications**

You see a list of the predefined applications supported. Select the predefined application, click the vertical ellipsis, and then click **Preview**. From the relevant Application pane, you can view the application services that can be discovered, and if the connections are identified, they will form an application. You can view the applications that are discovered by clicking **Operations › Applications › Applications Home**.

> **NOTE**
> If the same service instance is a member of more than one application (both known and custom), then the service is a part of the application that has the greatest number of services configured in the application definition.
>
> The discovered set of services that communicate with each other, should match at least 70% of the defined application. Only the matching ones are filtered out based on whether one service is defined on more than one application.

### Discovered Rule-Based Applications

VMware Aria Operations discovers applications based on user defined criteria such as object names, tags, or properties.

#### Where You View the Application Discovery Definitions

From the left menu, click **Operations › Configurations**. From the right pane, click the **Application Discovery: Rule Based** tile.

#### Application Discovery: Rule Based

The rule based application discovery provides you the following options.

- Specify grouping criteria for each rule, it can be tags, object names, or properties
- Apply additional filters to the rule to limit the scope of objects considered for grouping

You can discover applications based on the user defined criteria such as object names, tags, or properties. After the rules are created you can assign these rules to specific objects. VMware Aria Operations discovers applications based on these rules and displays them in Applications Home page.

#### Edit or Delete Application Rule

You can modify or delete the application definition. To edit or delete the application definition, click the vertical ellipsis next to the application and select **Edit** or **Delete**. If you select the edit option you are redirected to the **Edit Application Definition** page, here you can modify the application definition and save the changes. If you select the delete option, you are prompted with a confirmation message, click **Yes** to delete the application definition. Please note that if you delete the application definition, the applications that were discovered by that definition are also deleted.

### Alert for Service Unavailability

When a service is unavailable, an alert is triggered for the specific VM.

#### Alert for Service Unavailability

On a VM that is monitored, if one of the services is down, in the next collection cycle an alert is triggered.

| Alert Name | Symptom |
|---|---|
| One or more monitored service(s) are unavailable on the virtual machine. | Service is not available.<br>When the service is available again, the symptom disappears. |

The alert is canceled in the following scenarios:

- When all the discovered services are available again in the monitored VM.
- If the service is not available within 7 days.
- If you deactivate service monitoring for the monitored VM.

**Where You Find the Alert**

From the **Manage Services** page, ensure that the VM is monitored and one or more service(s) are unavailable on the VM. Select the VM, click **Show Details** to go to the summary page. Click **Alerts** from the toolbar, and then click the **Alerts** tab.

**Property for Service Unavailability**

You can view the property called Status for a service that has been discovered on the VM. For more information, see the topic called Services Properties.

**Service Discovery Metrics**

Service discovery discovers metrics for several objects. It also discovers CPU and memory metrics for discovered services.

**Virtual Machine Metrics**

Service Discovery discovers metrics for virtual machines.

**Table 101: Virtual Machine Metrics**

| Metric Name | Description |
|---|---|
| Guest OS Services\|Total Number of Services | Number of out-of-the-box and user-defined services discovered in the VM. |
| Guest OS Services\|Number of User Defined Services | Number of user-defined services discovered in the VM. |
| Guest OS Services\|Number of OOTB Services | Number of out-of-the-box services discovered in the VM. |
| Guest OS Services\|Number of Outgoing Connections | Number of outgoing connection counts from the discovered services. |
| Guest OS Services\|Number of Incoming Connections | Number of incoming connection counts to the discovered services. |

**Service Summary Metrics**

Service discovery discovers summary metrics for the service object. The object is a single service object.

**Table 102: Service Summary Metrics**

| Metric Name | Description |
|---|---|
| Summary\|Incoming Connections Count | Number of incoming connections. |
| Summary\|Outgoing Connections Count | Number of outgoing connections. |
| Summary\|Connections Count | Number of incoming and outgoing connections. |
| Summary\|Pid | Process ID. |

**Service Performance Metrics**

Service discovery discovers performance metrics for the service object. The object is a single service object.

**Table 103: Service Performance Metrics**

| Metric Name | Description |
|---|---|
| Performance metrics group\|CPU | CPU usage in percentage. |
| Performance metrics group\|Memory | Memory usage in KB. |
| Performance metrics group\|IO Read Throughput | IO read throughput in KBps. |
| Performance metrics group\|IO Write Throughput | IO write throughput in KBps. |

### Service Type Metrics

Service discovery discovers metrics for service type objects.

**Table 104: Service Type Metrics**

| Metric Name | Description |
|---|---|
| Number of instances | Number of instances of this service type. |

## Configuring Business Applications

A Business Application is a set of interconnected applications, services and hosts, which are configured to offer a service to the organisation. Business Applications can be internal, like organization email system or customer-facing, like an organization website. The Business Applications page in VMware Aria OperationsVMware Cloud Foundation Operations is where you see Business Applications and their health.

A Business Application is a container construct that represents a collection of interdependent hardware and software components that deliver a specific capability to support your business. You create a Business Application to determine how your environment is affected when one or more components in the Business Application experiences problems, and to monitor the overall health and performance of the Business Application. VMware Aria Operations can also display applications discovered through Service Discovery, and by importing applications from VMware Aria Operations for Networks.

VMware Cloud Foundation Operations can also display applications discovered through Service Discovery, and by importing applications from VMware Aria Operations for Networks SaaS and Ensemble.

### The Business Application Page

The Business Applications page is where you view all your Business Applications and their health in a sortable table. The health of the Business Application is determined by the aggregated health of the underlying objects. The **Health** column in the table displays the status.

#### Where You View the Business Applications Page

From the left menu, select **Operations › Configurations**. Under **Logical Groupings**, click the **Business Applications** tile to view the Business Applications page. All the available Business Applications and applications are displayed in the table.

The table displays the following types of objects:

**Table 105: Business Applications Page**

| Business Applications Source | Business Applications Type |
|---|---|
| Auto discovered and monitored applications which have more than one node | • Imported through VMware Aria Operations for Networks.<br>• Imported through VMware Aria Operations for Networks SaaS. |
| Imported from VMware Aria Hub. | |
| Manually created Business Applications. | • Can contain custom groups.<br>• Can contain any object to tier relationships.<br>• Can contain applications. |

**Columns in the Business Application Table**

The Business Applications table displays the following columns:

| Column Name | Description |
|---|---|
| Vertical Ellipses | Offers options for the user to perform actions on the Business Application.<br>You can also Enable Synthetic Monitoring, Disable Synthetic Monitoring, and Configure Synthetic Monitoring for the Business Application. |
| Health | Displays the health of the Business Application by showing different icons. On hovering the cursor over the icon, you can see the health percentage. |
| Name | Displays the name of the Business Application and an icon indicates the type of Business Application. |
| Description | Displays the description of the Business Application you entered when creating it. |
| Synthetic Monitoring State | Displays one of the following states for Synthetic Monitoring:<br>• Enabled<br>• Disabled<br>• Configured<br>• Not configured |
| Application Tags | Displays the application tags you entered when you created the Business Application. |
| Business Criticality | Displays one of the following criticality states for the Business Application:<br>• Medium<br>• Critical<br>• Low |
| Environment | Displays the environment where the Business Application resides in. |

## Actionable Options on the Business Applications Page

**Table 106: Actionable Options on the Business Applications Page**

| Option | Description |
|---|---|
| Vertical Ellipses | Click the vertical ellipses beside a Business Application to perform one of the following actions:<br>• Edit (Not available for VMware Cloud Foundation Operations created Business Applications)<br>• Delete<br>• Go to Details. This opens the Business Applications Summary tab in the Inventory page.<br>• Troubleshoot with Workbench<br>• Edit (Not available for VMware Cloud Foundation Operations created Business Applications)<br>• Delete<br>• Go to Details. This opens the Business Applications Summary tab in the Inventory page.<br>• Troubleshoot with Workbench<br>• Configure Synthetic Monitoring<br>• Enable Synthetic Monitoring<br>• Disable Synthetic Monitoring<br>• Edit Synthetic Monitoring |
| Filter | Filter the objects in the table by name. |
| Add Business Application | Click to create a new Business Application |
| Show Columns | Click to hide/show columns in the table. |
| Preview | Click on a Business Application name in the table to see a visual preview of the objects and their relationships. |

### Add Business Application

A Business Application is a first class citizen with tiers. Business Applications are groups of related objects in your environment that mimic an application in your business. You can create application to application relationships. Once you add a Business Application to VMware Aria OperationsVMware Cloud Foundation Operations, you can use the Business Applications page to track the health of objects in the application. You can optimize capacity, cost, and run What-if Analysis on Business Applications.

## How Business Applications Work

The tier is a convenient way to organize objects that perform a specific task in an application in your organization. For example, you can group all of your database servers together in a tier. Business Applications need not contain tier objects. Application or custom group objects can be connected to a Business Application without a tier object, and a tier object can contain no objects.

The objects in a tier are static. If the set of objects in a tier changes, you must manually edit the Business Application. However, if the Business Application is created by VMware Aria OperationsVMware Cloud Foundation Operations, the changes in the actual application will be reflected in the Business Application object without manual edit.

Construct a Business Application to view a particular segment of your business. The Business Application shows how the performance of one object affects other objects in the same Application, and helps you to locate the source of a problem. For example, if you have an application that includes all the database, Web, and network servers that process sales data

for your business, you see a yellow, orange, or red status if the application health is degrading. From the Business Applications page, you can investigate which server is causing or exhibiting the problem.

### The Add Business Application Screen

When you click **Add**, VMware Aria OperationsVMware Cloud Foundation Operations displays a blank canvas on the left pane, and options to select objects on the right pane. The **Select Members** section in the right pane has two tabs. The **Inventory** tab displays a list of all objects in you environment in a tree structure. You can search for objects here. The **Custom Groups** tab displays all the custom groups in your environment. You can use the search bar to search for custom groups.

Drag and drop objects from the right pane to the tier objects on the left pane. You can also drag and drop objects directly to the Business Application. Click **Save** to create the Business Application.

The new Business Application canvas is what you use to build your Business Application.

**Table 107: New Business Application Canvas**

| Option | Description |
|---|---|
| Vertical icon | Change the layout of the canvas from horizontal to vertical. |
| Undo icon | Undo the last action. |
| Redo icon | Redo the last action. |
| Fit View icon | Fit the Business Application structure within the boundary of the canvas. Use this option after you zoom into the canvas. |
| Add Tier icon | Add a new tier to the Business Application |
| Edit Tier icon | Available when you click a tier. Click the edit icon in a tier to edit the name and description of the tier. |
| Edit Business Application Details icon | Available when you click a Business Application. Click the edit icon in the Business Application to open the Business Applications Detail dialog box. Alternatively, click the DETAILS link next to the name of the Business Application. |
| Preview | Available for custom group and application objects when you have added objects to a tier. Click the preview icon after you have added an Application to a tier to see how services and objects within the application are related. This interactive element is displayed on the right pane. You can view links, switch to spherical view or list view. If you have too many objects to can use the filter option to find objects. |

The **Select Members** pane is where you select objects to add to the Business Application. Drag an object to add to a tier in the canvas. Applications and custom groups can be nested under a Business Application. Custom groups can be added to applications and vice versa.

To find an object, search by name. Each object listed includes identifier information to help distinguish between objects of similar names.

**Table 108: Select Members Pane**

| Option | Description |
|---|---|
| Search | Available in the **Inventory** and **Custom Groups** tabs. Search for objects to add to the Business Application. Filter |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | your search by clicking the filter icon. In the **Inventory** tab, by using the filter option, you can choose to see only objects which are Applications. |
| Inventory | Browse the inventory to select objects to add to the Business Application using drag and drop. |
| Custom Groups | Browse custom groups to add to the Business Application using drag and drop. You can nest custom groups or add them to the Business Application directly. You can add custom groups to Applications. |

The Business Applications Details dialog box is where you add information to build out a model of your Business Application resources through applications and other infrastructure objects in your environment.

**Table 109: Business Application Details**

| Option | Description |
|---|---|
| Description | Provide a description for your Business Application. |
| Application Tag | Provide a tag for your Business Application. |
| Business Criticality | This is your Business Application business criticality. Select between Medium, Critical and Low. Default is Medium. |
| Environment | This is where your Business Application is deployed. Select from one of the following options:<br>• DR<br>• Development<br>• Production<br>• Staging<br>• Test |

### Configure, Activate, or Deactivate Synthetic Monitoring

Synthetic monitoring can report API monitoring and check results in Prometheus metrics format which is pushed to VMware Cloud Foundation Operations for further analysis, visualization, and alerting purposes.

#### Where you find Synthetic Monitoring

From the left menu, click **Environment › Business Application**. On the Business Application page, click the vertical ellipses next to a Business Application name to configure, activate, or deactivate Synthetic Monitoring. Deatcivate Synthetic Monitoring appears only for a Business Application which has Synthetic Monitoring already activated. Activate Synthetic Monitoring appears only for a Business Application which has Synthetic Monitoring already configured.

#### Synthetic Monitoring Configuration Page

You can add up to 10 APIs for monitoring. Click the **ADD NEW API** button to add new APIs. The API configuration options appears in the bottom pane.

#### ADD NEW API Options

When you click the **ADD NEW API** button, you get the following configuration options:

| Option | Description |
|---|---|
| API Name | Provide a name for the API call. |
| API Request | Select the **HTTP Method** and provide the **Public Endpoint URL**. Options for the API type are:<br>• GET<br>• POST |
| Advanced Settings | Set number of API call retry attempts and timeout value. |
| Headers | In the **Headers** tab, select the content type of the body from the drop-down menu. The options are:<br>• application/json<br>• application/xml<br>Provide inputs for the **Key** and **Value** and click **TEST**. |
| Body | In the **Body** tab, provide the body for **API Request**. |
| Parameters | Enter **KEY** and **VALUE** properties |
| Authentication | In the **Authentication** tab, select a type of authentication for the API from the drop-down menu. The options are:<br>• None<br>• Basic Authentication<br>• Bearer Token<br>• OAuth Authentication<br>When you select a type of authentication apart from None, you get the options to provide credentials. The credentials depend on the authentication type.<br><br>Once you provide the authentication type, click **TEST** to check if VMware Cloud Foundation Operations can reach the API with the values you provided.<br><br>**NOTE**<br>When you upgrade to the latest version of VMware Aria OperationsVMware Cloud Foundation Operations, all credentials get unassigned. The VMware Aria OperationsVMware Cloud Foundation Operations administrator must assign the credentials from the **Orphan and Unassigned** page. For more information, see the topic, Managing Orphaned and Unassigned Content. |

Click **TEST** to test the API request. Click **Save** after you provide these values..

VMware Cloud Foundation Operations monitors the APIs at a frequency displayed in the **Frequency** property in the top pane after you activate Synthetic Monitoring for the Business Application. Each Business Application can have up to 10 API requests registered. VMware Cloud Foundation Operations can have up to 20 Business Applications with Synthetic Monitoring activated.

You can view the results of synthetic monitoring in the Synthetic Monitoring tab. See, the topic, Synthetic Monitoring Tab.

# Application Integration

You can integrate Application Performance Monitoring tools to discover applications in VMware Cloud Foundation Operations.

The integrations for Application Performance Monitoring tools are enabled using the following applications.

- AppDynamics. For product documentation, see VMware Aria Operations Management Pack for AppDynamics.
- New Relic. For product documentation, see VMware Aria Operations Management Pack for New Relic.
- DataDog. For product documentation, see VMware Aria Operations Management Pack for Datadog.
- Dynatrace. For product documentation, see VMware Aria Operations Management Pack for Dynatrace.

### Application Discovery

Integrating the Application Performance Monitoring tools to discover applications, enhances the troubleshooting abilities by retrieving the application topology and key metrics.

#### How Can I View Applications Available for Integration

From the left menu, click **Environment › Applications › Application Integration**. Click **Discover Applications** to view the applications available for integration and the applications that are successfully integrated. For details, see Integrating Applications.

#### Where Can I Find the Discovered Applications

Once the application is integrated, the discovered applications are displayed in the **Data Sources › Integrations** page after a few collection cycles.

**Table 110: Application Integration Page Options**

| Options | Descriptions |
|---|---|
| All Filters | You can filter the discovered applications by Name, Object Type, and Adapter Type. |
| Application Name | Displays the name of the application. Click this link to view the details in the **Summary** page. |
| Object Type | Displays the type of object. |
| Source | Displays the source of the application. |
| Adapter Type | Displays the name of the adapter. |
| Active Alerts | Displays the active alerts associated with the application. |

### Integrating Applications

You can integrate Application Performance Monitoring tools to discover applications in VMware Cloud Foundation Operations. These integrations enhance the troubleshooting abilities by retrieving the application topology and key metrics.

The **Application Integration** page displays the applications available for integration and the applications that are successfully integrated. The integrations for Application Performance Monitoring tools are enabled using the following applications.

- AppDynamics
- New Relic
- DataDog
- Dynatrace

**How Can I Download Applications for Integration**

From the left menu, click **Environment › Applications › Application Integration**.

Click **Configure Now** and in the **Integrations › Account Types** page, click the required management pack to automatically download it. For details, see The Integrations Page.

You can also download the management pack from the **Repository** tab under **Data Sources › Integrations**. For details, see The Repository Tab.

**How Can I Add Accounts for Solutions**

From the left menu, click **Data Sources › Integrations › Accounts**. Click **Add Account**, and then select the solution you want to manage.

You can also add accounts from the **Repository** tab under **Data Sources › Integrations**. For details, see The Repository Tab.

## Configuring Ping Adapter Instances

In VMware Aria OperationsVMware Cloud Foundation Operations, you can configure the Ping functionality to verify the availability of end points that exist in your virtual environment. The ping functionality is configured at the adapter instance for IP addresses, group of IP addresses, and FQDN.

- If you have multiple adapter instances running on different collectors and both are pinging the same address, you can still get statistics from both the adapter instances for the same IP.
- The FQDN names are checked for validity, the FQDN validation relies on RFC1034 and RFC1123, and only top level domains of the internet are validated. The `.local` domain is not supported as it does not fall into the list of top-level domains in the Domain Name System (DNS) of the Internet.

1. From the left menu, click **Administration › Integrations**.
2. In the Accounts tab, click **Add Account**
3. Click **Other** to filter the list of accounts. The Ping account tile is displayed after filtering out other tiles.
4. Click the Ping adapter instance.
5. Click **Yes** in the diaglog box in the dialog box which opens. This will install the management pack.
6. Configure the Ping adapter instance.

| Option | Description |
|---|---|
| Name | Enter a name for the adapter instance. |
| Description | Enter the description of the adapter instance. |
| Unique Name | Specify the name for the adapter instance. You can use the name to view the metrics published for the adapter instance. |
| Address List | Specify the IP address, IP address range, and the FQDN which must be pinged. |
| Configuration Filename | Specify the name of the configuration file. The configuration file contains the IP addresses, CIDR information, and FQDN details as a comma-separated file. |
| Collectors/Groups | Select the collector from which this adapter instance must run. |
| Validate Connection | Click to check whether the connection is successful or not. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Advanced Settings | To configure the advanced settings, click the drop-down menu. |
| Wait Interval Time (second) | Specify the time interval in seconds to wait before running the next batch. Range: 0-300 seconds. |
| Batch Size | Specify the number of request packets to send to each target. Range: 20-100. |
| Interval (millisecond) | Specify the time the fping waits between successive packets to an individual targets. Greater or equal to 2000 milli seconds. |
| DNS Name Resolve Interval | Specify the time at which you must resolve the DNS name for the next cycle. Minimum value is 15 minutes. |
| Packet Size | Specify the byte size of the packet when you ping. Range: 56-65536 bytes. |
| Don't Fragment | Select **False** to fragment the packet and **True** to not fragment the packet. |
| Generate FQDN Child IPs | Select **True** to create IP objects by resolved names and add as child of FQDN. |

7. Click **ADD**.

After you configure the Ping adapter instance, you can view the adapter details from **Administration › Integrations › Repository**.

## VMware Aria Automation

VMware Aria Automation extends operational management capabilities of the VMware Aria OperationsVMware Cloud Foundation Operations platform to provide the cloud aware operational visibility of the cloud infrastructure. VMware Aria Automation helps you to monitor the health, efficiency, and capacity risks associated with the imported cloud accounts.

You can use VMware Aria Automation to perform some of the following key tasks:

- Gain visibility into the performance and health of cloud zones integrated with VMware Aria OperationsVMware Cloud Foundation Operations.
- Import and synchronize existing cloud accounts from VMware Aria Automation to VMware Aria OperationsVMware Cloud Foundation Operations.
- Manage the workload placement of VMs that are part of the clusters managed by VMware Aria Automation.
- Integrate and troubleshoot vSphere endpoint issues associated with VMware Aria Automation using the VMware Aria OperationsVMware Cloud Foundation Operations dashboard.

**NOTE**
In this release we support only vSphere endpoints.

### Advanced Workload Placement for Allocation Model

VMware Aria OperationsVMware Cloud Foundation Operations supports allocation aware advanced workload placement between VMware Aria Automation and VMware Aria OperationsVMware Cloud Foundation Operations for initial placement of virtual machines. The placement recommendation is based on allocation settings that are defined in the VMware Aria OperationsVMware Cloud Foundation Operations policy. The advanced workload placement awareness for allocation works with the existing demand model.

## Need for Allocation Model Awareness

The advanced workload placement based on demand model, depends on the actual demand for resources in a cluster and datastore. If the allocation is based only on the utilization of resources, then it might result in over allocation or over provisioning of resources in clusters. To avoid this, VMware Aria OperationsVMware Cloud Foundation Operations provides allocation model awareness. Allocation model awareness addresses issues related to over allocation or over provisioning by providing you the option of setting the appropriate overcommit ratios in the VMware Aria OperationsVMware Cloud Foundation Operations policy.

The advanced workload placement feature uses the demand model by default and cannot be turned off. To activate the allocation model in addition to the demand model, in VMware Aria OperationsVMware Cloud Foundation Operations configure the appropriate overcommit ratios in policy settings for the preferred clusters, datastores, and datastore clusters.

## Prerequisites

- VMware Aria OperationsVMware Cloud Foundation Operations is configured as an endpoint in VMware Aria Automation,this happens automatically in case of VMware Aria OperationsVMware Cloud Foundation Operations
- Advanced placement policy in the Cloud zone must be activated
- vCenter Cloud account instance on which the initial provision is done should be same across VMware Aria OperationsVMware Cloud Foundation Operations and VMware Aria Automation
- Overcommit ratios in VMware Aria OperationsVMware Cloud Foundation Operations policy must be configured for the following:
  - For CPU and Memory overcommit ratio - Cluster Compute Resource
  - For Disk overcommit ratio - Datastore and Datastore Cluster

## How to Activate Advanced Placement Policy

To activate the advanced placement policy, in VMware Aria Automation Cloud Assembly Service, navigate to **Infrastructure** > **Cloud Zone** > **Summary** tab and set the placement policy to **ADVANCED**.

> **NOTE**
> If VMware Aria OperationsVMware Cloud Foundation Operations returns no recommendations, then under the placement policy, you can specify if you want VMware Aria Automation to fall back to its default placement using the toggle option.

## How to Activate Allocation Awareness in Advanced Workload Placement Feature

To activate Allocation awareness in VMware Aria OperationsVMware Cloud Foundation Operations perform the following actions.

1. From the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile.
2. Select the Active policy which is assigned to the Cluster Compute resource under Cloud Zones and make the required changes.
3. Click **Edit Policy** and navigate to the **Capacity** tile.
4. Select the **Object Type** as Cluster Compute Resource and activate the Allocation Model.
5. Set the overcommit ratio as per your requirement and click **Save**.
6. Repeat Step 3 to Step 5 for datastore cluster, vSAN datastores and local datastores.

After the configuration is completed, VMware Aria Automation Cloud makes a provisioning request to VMware Aria OperationsVMware Cloud Foundation Operations, and the advanced workload placement engine calculates the recommendation and shares it with VMware Aria Automation. To know the changes after the configuration, view the Allocation metrics and Demand metrics for the cluster.

**NOTE**
The allocation model awareness enhancement is limited to advanced workload placement feature. This capability is not extended to workload optimization feature.

### Integrate VMware Aria Automation SaaS Service with VMware Cloud Foundation Operations SaaS Service

The integration of the VMware Cloud Foundation Operations service and VMware Aria Automation Cloud service happens automatically if your organization has access to both the services. If you add VMware Aria Automation Cloud services first and then add VMware Cloud Foundation Operations cloud services, VMware Cloud Foundation Operations autoconfigures the VMware Aria Automation Cloud accounts.

#### How Cloud Integration Works

Once the integration of VMware Aria Automation cloud service with VMware Cloud Foundation Operations is complete. VMware Cloud Foundation Operations service provides the following information about the integrated VMware Aria Automation cloud account:

- Import cloud accounts defined in VMware Aria Automation Cloud to VMware Cloud Foundation Operations.
- View cloud zones defined in VMware Aria Automation cloud in VMware Cloud Foundation Operations.
- Modifications done to Cloud Zones in VMware Aria Automation cloud are reflected in VMware Cloud Foundation Operations.
- View the integrated cloud adapter instance in VMware Aria Automation inventory objects.
- List the objects related to the integrated cloud adapter in VMware Cloud Foundation Operations inventory.
- Access to the objects in your VMware Aria Automation cloud account is based on your user role. An organization administrator has access to all the objects in your cloud environment, but an organization member has limited access to the objects in your cloud environment.
- Updates done to the objects or data of VMware Aria Automation cloud are reflected in VMware Cloud Foundation Operations as well.
- After the VMware Aria Automation cloud and VMware Cloud Foundation Operations integration is complete, all the dashboards of VMware Aria Automation are created in VMware Cloud Foundation Operations.

You can perform reset and upgrade operation on VMware Aria Automation cloud accounts. If you reset the VMware Aria Automation account, you create an environment and all the data related to the old setup is removed from the system.

If you upgrade the VMware Aria Automation cloud account, the historical data is retained in the VMware Cloud Foundation Operations account.

#### Importing Accounts from VMware Aria Automation

You can import and synchronize existing cloud accounts from VMware Aria Automation to VMware Aria OperationsVMware Cloud Foundation Operations. Click **Import Accounts from VRA** > **Import Accounts** to list all the cloud accounts associated with vCenter and Microsoft Azure that are not managed by VMware Aria OperationsVMware Cloud Foundation Operations. You can select and import these accounts into VMware Aria OperationsVMware Cloud Foundation Operations directly with existing credentials as defined in VMware Aria Automation or add or edit the credentials before the import process. The **Import Accounts from VRA** option is hidden from the user until the integration with VMware Aria Automation is activated from the integration page under **Administration › Integrations › Accounts** or **Repository** tabs.

- Verify that VMware Aria Automation is activated from **Administrations › Integrations › Accounts** in VMware Aria OperationsVMware Cloud Foundation Operations.
- Verify that you know the vCenter credentials that have sufficient privileges to connect and collect data.
- Verify that the user has privileges of Organizational Owner and Cloud Assembly administrator set in VMware Aria Automation.

1. From the left menu, go to **Administration › Integrations › Accounts** tab, click on the horizontal ellipses, and then select **Import Accounts from VRA**.

2. From the **Import Accounts** page, select the cloud account you want to import.
3. To override an existing credential from VMware Aria Automation.
   – Select the existing credential from the **Credential** drop-down menu and click **Save**.
   – To add a new credential, click the plus icon next to the **Credential** drop-down menu and enter the credential details and click **Save**.
4. Select the collector/group from the drop-down menu.
5. Click **Validate** to verify that the connection is successful.
6. Click **Import**.

The imported cloud account is listed in the **Administration › Integrations › Accounts** page. After the data collection for the cloud account is complete the configuration status changes from **Warning** to **OK**.

### Supported VMware Aria Automation Versions

VMware Aria Automation is supported on vRealize Operations 8.6 version. Workload placement for day 1 operations is supported from vRealize Automation 7.3 onwards with vRealize Operations 6.6 and above. Workload placement for day 2 operations is supported from vRealize Automation 7.5 onwards with vRealize Operations 7.0 and above.

### Object Types

VMware Aria Automation brings in cloud accounts and their relationships from VMware Aria Automation into VMware Aria OperationsVMware Cloud Foundation Operations for operational analysis. You can use the following items in the virtual infrastructure as object types in VMware Aria OperationsVMware Cloud Foundation Operations.

- Cloud Zone
- Blueprint
- Project
- Deployment
- Cloud Account
- User
- Organization
- Cloud Automation Services World

#### Workload Placement

In VMware Aria OperationsVMware Cloud Foundation Operations, you can configure VMware Aria Automation instances to work with VMware Aria OperationsVMware Cloud Foundation Operations instances. Using VMware Aria OperationsVMware Cloud Foundation Operations you can monitor the placement of existing workloads and optimize the resource usage.

- Verify that the user has privileges of Organizational Owner and Cloud Assembly Administrator set in VMware Aria Automation.
- You must know the vCenter credentials and have the necessary permissions to connect and collect data.
- Verify that VMware Aria Automation is activated from **Administration** > **Integrations** in VMware Aria OperationsVMware Cloud Foundation Operations. For more information, see Configuring VMware VMware Aria Automation with VMware Aria Operations.
- VMware Aria OperationsVMware Cloud Foundation Operations must have the same vCenter Cloud Account configured to match with VMware Aria Automation.
- Ensure that integration is activated for VMware Aria OperationsVMware Cloud Foundation Operations and VMware Aria Automation.

1. From the left menu, click **Capacity** > **Workload Placement**.
2. Click the **View** filter drop-down menu and select the **VRA Managed** objects.
   All the Cloud Zones related to the vCenter are displayed in VMware Aria OperationsVMware Cloud Foundation Operations.
3. Click the **Cloud Zone** you want to optimize.
4. Based on the operational intent, click **Optimize Now**.

   The system creates an optimization plan, which depicts BEFORE and (projected) AFTER workload statistics for the optimization action.

5. If you are satisfied with the projected results of the optimization action, click **NEXT**.
6. Review the optimization moves, then click **BEGIN ACTION**.

   In the scope of VMware Aria Automation integration, VMware Aria OperationsVMware Cloud Foundation Operations sends a move migration request directly to VMware Aria Automation. In the earlier versions, the migration request was sent to the vCenter.

To verify that the optimization action is complete, select **Administration** from the left menu, and click **Recent Tasks** . In the **Recent Tasks** page, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, first filter on Starting Time and scroll to the time when you began the action, then select the Object Name filter. Finally, enter the name of one of the VMs in the rebalance plan.

### Pricing for VMware Aria Automation Components in VMware Aria OperationsVMware Cloud Foundation Operations

After you integrate VMware Aria Automation private cloud adapter instances with VMware Aria OperationsVMware Cloud Foundation Operations, you can calculate the cost of deployments, projects, and virtual machines of the selected cloud adapter. Pricing provides an overview of the costs related to the cloud environment, cloud resources, and the costs associated with the project.

#### How the Pricing Works in VMware Aria Automation

- VMware Aria OperationsVMware Cloud Foundation Operations understands the constructs defined in VMware Aria Automation and calculates the CPU, RAM, Storage and Additional prices for Projects, Deployments, and virtual machines.
- A single project can have multiple deployments and a single deployment can have multiple virtual machines associated with the deployment.
- Pricing for multiple virtual machines associated with the deployment is the sum of all the resources associated with individual virtual machines.
- If a single project has multiple deployments, then the project pricing is equal to the sum of individual deployments. The deployment can have multiple virtual machines and resources associated with it.
- On day one, the pricing is equal to the cost of resources defined in VMware Aria OperationsVMware Cloud Foundation Operations.
- On day two, the price is calculated using the following formula.
  - Cost of resources for the present day – Cost of resources for the previous day
- If in case the pricing does not happen as per the definition, then the partial price is set to true, and the pricing is calculated based on the previous days price.
- In VMware Aria OperationsVMware Cloud Foundation Operations, the following new dashboards are included to view the pricing details for the VMware Aria Automation instances.
  - Cloud Automation Environment Overview
  - Cloud Automation Project Cost Overview

  - Cloud Automation Resource Consumption Overview
  - Cloud Automation Top-N Dashboard

**Data Collection Enhancements in VMware Aria Automation for Pricing in VMware Aria OperationsVMware Cloud Foundation Operations**

The following enhancements have been made for the data collection process from VMware Aria Automation for pricing purposes.

- Collect cloud zones with relation to clusters and resources pools from VMware Aria Automation to VMware Aria OperationsVMware Cloud Foundation Operations.
- Collect Projects from VMware Aria Automation with relation to deployments.
- Include project, cloud zone, and blueprint as properties in virtual machines that are deployed in VMware Aria Automation.

**Upfront Price Support for VMware Aria Automation Private Cloud Components**

VMware Aria OperationsVMware Cloud Foundation Operations supports upfront pricing for VMware Aria Automation 8.x in the following ways:

- VMware Aria OperationsVMware Cloud Foundation Operations uses rate cards to provide upfront cost estimates of catalog items just before deployment.
- VMware Aria Automation retrieves the deployment cost and estimated cost from VMware Aria OperationsVMware Cloud Foundation Operations.
- VMware Aria Automation user interface allows you to customize the pricing policies and assign them to the projects or cloud zones.
- If VMware Aria Automation does not specify the pricing policy, then the price is calculated using the VMware Aria OperationsVMware Cloud Foundation Operations cost calculation policy.
- If a custom pricing policy is set for a price calculation, then the deployment and upfront catalog price computation is done as per the custom policy.

**Upfront Price Support for VMware Cloud on AWS Resources**

VMware Aria OperationsVMware Cloud Foundation Operations supports upfront pricing for VMware Cloud on AWS resources in the following ways:

- VMware Aria OperationsVMware Cloud Foundation Operations supports upfront pricing for VMware Cloud on AWS only if rate-based pricing is configured in VMware Aria Automation for VMware Cloud on AWS resources.
- VMware Aria OperationsVMware Cloud Foundation Operations does not support cost-based computation for VMware Cloud on AWS resources.

**Upfront Pricing Support for Metering Policy**

VMware Aria OperationsVMware Cloud Foundation Operations supports upfront pricing for metering policy in the following ways:

- VMware Aria OperationsVMware Cloud Foundation Operations supports tag-based metering policy for cost calculation using virtual machines with specific key and value. Virtual machines can be charged on a per day basis.
- VMware Aria OperationsVMware Cloud Foundation Operations supports metering policy with one-time charges for virtual machines on a daily basis.
- VMware Aria OperationsVMware Cloud Foundation Operations supports metering policy for specific operating systems.
- VMware Aria OperationsVMware Cloud Foundation Operations supports custom properties in metering policy for calculating cost of resources in virtual machines.

**Configuring VMware Aria Automation with VMware Aria OperationsVMware Cloud Foundation Operations**

To access VMware Aria Automation instance and troubleshoot automation issues using VMware Aria OperationsVMware Cloud Foundation Operations, you must configure the VMware Aria Automation adapter in VMware Aria OperationsVMware Cloud Foundation Operations.

- Verify that you know the FQDN/IP address, user name, and password of the VMware Aria Automation instance you have installed.
- Ensure that the VMware Aria Automation user has both organizational owner and Cloud Assembly administrator permissions.
- When you use a domain account for configuration, you must specify the username to validate your authentication. If you try to validate using the `<domain>\<username>` and `<username>@<domain>` formats the authentication might not work.
- vRealize Operations 8.2 or later supports one-to-one integrations with VMware Aria Automation, you can integrate one instance of vRealize Operations 8.2 or later with one instance of VMware Aria Automation.
- VMware Aria Automation or later supports one-to-many integration with vRealize Operations 8.2 or later, you integrate more than one vRealize Operations 8.x instance with one VMware Aria Automation end point.
- To know more about integration between VMware Aria Automation and VMware Aria Operations, refer to the section Integrating with VMware Aria Operations in *VMware Aria Automation Product Documentation.*

1. From the left menu, select **Administration › Integrations**.
2. From the **Accounts** tab in the **Integrations** page, click **Add Account** and then select the VMware VMware Aria Automation card. You can also activate the management pack from the **Repository** tab, where you will find the card in the Available Integrations section. After the management pack is activated, click **Add Account**.
3. In the VMware Aria Automation page, enter the FQDN or IP address of the VMware Aria Automation instance to which you want to connect.
4. Set **Auto Discovery** to true.
5. To add credentials, click the plus sign.
   a) In the Credential name text box, enter the name by which you are identifying the configured credentials.
   b) Enter the user name and password for the VMware VMware Aria Automation instance.
   c) Click **OK**.
   You have configured credentials to connect to a VMware VMware Aria Automation instance.
6. From the **Collectors/Groups** drop-down menu, select the collector group.
7. Click **Validate Connection** to verify that the connection is successful.
8. Review and accept the Server certificate.
9. Click **Advanced Settings**.
10. From the **User Count** drop-down menu, select the number of user resources to be imported from VMware Aria Automation.

    The User Count options are 20, 100, 200, 300, 400, and All Users.

11. Click **Save** to save the adapter instance.

After integrating the VMware Aria Automation adapter instance with VMware Aria OperationsVMware Cloud Foundation Operations, you can view the VMware Aria Automation adapter data from the VMware Aria OperationsVMware Cloud Foundation Operations dashboard.

### Support for VMware Aria Automation Management Pack Cloud Services in VMware Aria Operations SaaS

The VMware Aria OperationsVMware Cloud Foundation Operations extends operational management capabilities to Cloud Automation Services Management Pack, using VMware Aria OperationsVMware Cloud Foundation Operations you can retrieve cloud accounts, cloud zones, projects, blueprints, deployment, and virtual machines associated with VMware Aria Automation.

Using the VMware Aria Automation Management Pack for Cloud Services you can perform the following tasks in your cloud environment:

- Integrate VMware Aria Automation Management Pack cloud services with the VMware Aria OperationsVMware Cloud Foundation Operations at the organization level.

- Integrate VMware Aria OperationsVMware Cloud Foundation Operations specific workload placement engine with vRealize Automation 8.x workload provisioning and management engine for the optimal placement of resources.
- View Cloud Automation dashboards to monitor and troubleshoot objects in your cloud infrastructure.
- Verify that the existing cloud accounts from VMware Aria Automation are imported to VMware Aria OperationsVMware Cloud Foundation Operations.
- View the inventory details of VMware Aria Automation objects discovered in VMware Cloud Foundation Operations.
- Retrieve cloud zones defined in VMware Cloud Automation Services (CAS) into VMware Aria OperationsVMware Cloud Foundation Operations.

**Managing Public Cloud Endpoints with VMware Aria Automation Integration**

With the VMware Aria Operations management pack, you can monitor deployments made to public cloud endpoints such as Amazon Web Services (AWS), Google Cloud Platform, and Microsoft Azure. You can monitor the performance, health, utilization, and availability attributes of deployments made to public cloud endpoints.

This integration supports the VMware Aria Operations management packs for public clouds – MP for AWS, GCP, and Azure. Hence, the public cloud MPs are a prerequisite for this enhancement.

VMware Aria Operations will display the VMware Aria Automation deployed by public cloud resources as long as the same resources are also monitored by VMware Aria Operations for the respective public cloud MP.

This enhancement also shows all the cloud accounts and cloud zones that are part of public cloud end points.

> **NOTE**
> To monitor the VMware Aria Automation resources deployed to AWS, GCP, and Azure, you must ensure that the cloud accounts are configured in both VMware Aria Automation and VMware Aria Operations.

**Import Account Functionality**

Users may import AWS, Azure, and vCenter accounts automatically through the **Import Account** option. However, for the GCP account you should manually add the cloud account through **Add Account** option. The credential used to configure the adapter must match the one used in VMware Aria Automation. The traversal specification of VMware Aria Automation has been enhanced to include details about AWS, Azure, and GCP accounts.

**Cloud Zones in VMware Aria OperationsVMware Cloud Foundation Operations**

Cloud zones help you to group a set of compute resources and assign capability tags to the zone. The cloud zone is based on accounts/regions, so you must have at least one cloud account configured before you can create a cloud zone. Cloud zones define where and how blueprints configure deployments. You can have one or many cloud zones assigned to each project based on priority and limits.

**How Cloud Zones Work**

After you integrate VMware Aria Automation with VMware Aria OperationsVMware Cloud Foundation Operations, you can retrieve cloud zones into VMware Aria OperationsVMware Cloud Foundation Operations. The **Cloud Zones** option is hidden from the user until the integration with VMware Aria Automation is activated from the integration page under **Administration** > **Integrations**.

The Cloud Zones option is activated in VMware Aria OperationsVMware Cloud Foundation Operations, only if the following conditions are met.

- VMware Aria Automation instance is integrated successfully in VMware Aria OperationsVMware Cloud Foundation Operations**Administration** > **Integrations**.
- VMware Aria Automation objects are discovered in VMware Aria OperationsVMware Cloud Foundation Operations.
- VMware Aria Automation accounts and Aria Operations vCenter Cloud Accounts are synchronized.

All the Cloud Zone objects which are existing in VMware Aria Automation environment, are discovered in VMware Aria OperationsVMware Cloud Foundation Operations. Cloud zones, whose dependent clusters are not discovered in VMware Aria OperationsVMware Cloud Foundation Operations, are not represented in Capacity Overview, Reclaim, and Workload Optimization pages.

**Cloud Zones List**

You can view the list of cloud zones that exist in your environment. In this view, you can click a cloud zone to display all the resources and objects that are associated with the cloud account. When you click the Cloud Zone, you are directed to the standard object summary page of the cloud account.

**Where You Find Cloud Zones**

Select **Environment** in the menu and click **Cloud Zones** tab.

**Cloud Zone Tab Options**

| Option | Description |
|---|---|
| Name | Displays the name of the selected cloud zone. |
| Cloud Account | Displays the cloud accounts associated with the cloud zone. |
| Resources | Displays the cloud account resources associated with the cloud zone.<br>**NOTE**<br>If the resource field is empty, it means VMware Aria OperationsVMware Cloud Foundation Operations does not have a corresponding vCenter Cloud Account for that associated Cloud Zone. Add a new vCenter Cloud Account manually or use the Import Cloud Account option from the Cloud Account page. |
| Capability Tags | Displays the capability tags associated wit the cloud zone. |

## vSAN

You can make vSAN operational in a production environment by using dashboards to evaluate, manage, and optimize the performance of vSAN objects and vSAN-activated objects in your vCenter system.

vSAN extends the following features:

- Discovers vSAN disk groups in a vSAN datastore.
- Identifies the vSAN-activated cluster compute resource, host system, and datastore objects in a vCenter system.
- Automatically adds related vCenter components that are in the monitoring state.
- Support for vSAN datastores in workload optimization with cross-cluster rebalance actions.
  - You can move VMs from one vSAN datastore to another vSAN datastore.
  - You can optimize the container if all the vSAN clusters are not in resync state.
  - VMs with different storage policies for each disk or VMs with different types of storage for each disk will not be moved.
  - You can generate a rebalance plan only if sufficient disk space is available at the destination vSAN datastore (The vSAN datastore slack space will also be considered).

- The storage policy assigned to the VM will be considered during the workload optimization (Compatibility check is performed against the storage policy).
- VM migration from vSAN datastore to vSAN stretched clusters is not supported.

**Configure a vSAN Adapter Instance**

When configuring an adapter instance for vSAN, you add credentials for a vCenter. In the earlier versions of VMware Aria OperationsVMware Cloud Foundation Operations, the vSAN solution was installed as part of the VMware Aria OperationsVMware Cloud Foundation Operations installation. Now, in case of a new installation the vSAN solution is pre-bundled as part of VMware Aria OperationsVMware Cloud Foundation Operations OVF, you must install the vSAN solution separately.

Only vCenter systems that are configured for both the vCenter adapter and the vSAN adapter appear in the inventory tree under the vSAN and Storage Devices. Verify that the vCenter that you use to configure the vSAN adapter instance is also configured as a vCenter adapter instance for the VMware vSphere® solution. If not, add a vCenter adapter instance for that vCenter.

You must open port 5989 between the host and any VMware Aria OperationsVMware Cloud Foundation Operations node on which the vSAN adapter resides. This is applicable when the vSAN version in vSphere is 6.6 or lower.

You must have a vCenter Adapter instance configured and monitoring the same vCenter Server that is used to monitor the vSAN and Storage Devices.

To know how to install the Native Management Packs, see The Repository Tab.

1. From the left menu, select **Administration** and then click **Integrations › ADD**.
2. From the **Account Types** page, select the vCenter instance from the list and then click the **vSAN** tab.
3. To use the vCenter for activating vSAN, move the **vSAN configuration** option to the right.

   **NOTE**
   Once the vSAN adapter instance is activated and saved, the activated vSAN configuration option is not visible.

4. The credentials provided for the vCenter instance are also used for vSAN adapter instance. If you do not want to use these credentials, you can click **Use alternate credentials** option.
   a) Click the plus sign next to the Credential field and enter the details in the **Manage Credentials** dialog box.
   b) Enter the credential name, vCenter user name, and password and click **OK**.
5. Choose **Enable SMART data collection**, to activate SMART data collection for physical disk devices.
6. Click **Add**.
   The vSAN configuration is activated for the cloud account.
7. Click **Test Connection** to validate the connection with your vCenter instance.
8. Accept the vCenter security certificate.
9. Click **Save Settings**.

The adapter is added to the Adapter Instance list and is active.

To verify that the adapter is configured and collecting data from vSAN objects, wait a few collection cycles, then view application-related data.

- Inventory. Verify that all the objects related to the vSAN instance are listed. Objects should be in the collecting state and receiving data.
- Dashboards. Verify that vSAN Capacity Overview, Migrate to vSAN, vSAN Operations Overview, and Troubleshoot vSAN, are added to the default dashboards.

- Under **Inventory › vSAN and Storage Devices**, verify that the vSAN hierarchy includes the following related vCenter system objects:
  - vSAN World
  - Cache Disk
  - Capacity Disk
  - vSAN-activated vCenter clusters
  - vSAN Fault Domains (optional)
  - vSAN-activated Hosts
  - vSAN Data stores
  - vSAN Disk Groups
  - vSAN Data store related VMs
  - vSAN Witness Hosts (optional)

### Verify that the Adapter Instance is Connected and Collecting Data

You configured an adapter instance of vSAN with credentials for a vCenter. Now you want to verify that your adapter instance can retrieve information from vSAN objects in your environment.

To view the object types, from the left menu, click **Inventory › Adapter Instances › vSAN Adapter Instance › <*User_Created_Instance*>**.

**Table 111: Object Types that vSAN Discovers**

| Object Type | Description |
|---|---|
| vSAN Adapter Instance | The VMware Aria Operations Management Pack for vSAN instance. |
| vSAN Cluster | vSAN clusters in your data center. |
| vSAN Datastore | vSAN datastores in your data center. |
| vSAN Disk Group | A collection of SSDs and magnetic disks used by vSAN. |
| vSAN Fault Domain | A tag for a fault domain in your data center. |
| vSAN Host | vSAN hosts in your data center. |
| vSAN Witness Host | A tag for a witness host of a stretched cluster, if the stretched cluster feature is activated on the vSAN cluster. |
| vSAN World | A vSAN World is a group parent resource for all vSAN adapter instances. vSAN World displays aggregated data of all adapter instances and a single root object of the entire vSAN hierarchy. |
| Cache Disk | A local physical device on a host used for storing VM files in vSAN. |
| Capacity Disk | A local physical device on a host used for read or write caching in vSAN |

The vSAN adapter also monitors the following objects discovered by the VMware vSphere adapter.

- Cluster Compute Resources
- Host System
- Datastore

1. In the menu, click **Administration** and then in the left pane, click **Configuration › Inventory** .
2. In the list of tags, expand **Adapter Instances** and expand **vSAN Adapter Instance**.
3. Select the adapter instance name to display the list of objects discovered by your adapter instance.
4. Slide the display bar to the right to view the object status.

| Object Status | Description |
|---|---|
| Collection State | If green, the object is connected. |
| Collection Status | If green, the adapter is retrieving data from the object. |

5. Deselect the adapter instance name and expand the **Object Types** tag.
   Each Object Type name appears with the number of objects of that type in your environment.

If objects are missing or not transmitting data, check to confirm that the object is connected. Then check for related alerts.

To ensure that the vSAN adapter can collect all performance data, the Virtual SAN performance service must be activated in vSphere. For instructions on how to activate the service, see Turn on Virtual SAN Performance Service in the VMware Virtual SAN documentation.

If the Virtual SAN performance service is deactivated or experiencing issues, an alert is triggered for the vSAN adapter instance and the following errors appear in the adapter logs.

```
ERROR com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics

   - Failed to collect performance metrics for Disk Group

com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics

   - vSAN Performance Service might be turned OFF.

com.vmware.adapter3.vsan.metricloader.VsanDiskgroupMetricLoader.collectMetrics

   - (vim.fault.NotFound)

   {

   faultCause = null,

   faultMessage = (vmodl.LocalizableMessage)

      [

         com.vmware.vim.binding.impl.vmodl.LocalizableMessageImpl@98e1294

      ]

   }
```

### vSAN Log Analytics Enhancements

When VMware Aria OperationsVMware Cloud Foundation Operations is integrated with VMware Cloud Foundation Operations for logs, you can view and troubleshoot VMware Cloud Foundation Operations for logs object issues within VMware Aria OperationsVMware Cloud Foundation Operations. Earlier you could troubleshoot issues related only to vCenter objects, but now you can troubleshoot issues related to vSAN also.

The enhancements to vSAN log analytics include use of specific queries to retrieve log information for the following vSAN objects:

- vSAN Cluster
- Witness Host
- Disk Group
- Cache Disk

- Capacity Disk

**Where You Find vSAN Object Logs**

Navigate to the vSAN Object Details page, and click the **Logs** tab.

> **NOTE**
> If you are not logged in to VMware Cloud Foundation Operations for logs, then VMware Aria OperationsVMware Cloud Foundation Operations prompts you to log in to VMware Cloud Foundation Operations for logs with your login credentials.

VMware Aria OperationsVMware Cloud Foundation Operations uses special queries for each object type. Using the special queries for vSAN objects, you can perform the following actions:

- View interactive analytics for the selected vSAN object.
- Retrieve log details for the vSAN object.
- Analyze and troubleshoot issues related to the vSAN object.

## VMware Aria Operations for Networks

The VMware Aria Operations for Networks adapter activates integration of VMware Aria OperationsVMware Cloud Foundation Operations with VMware Aria Operations for Networks. VMware Aria Operations for Networks provides network visibility and analytics to minimize risk during application migration, optimize network performance, manage and scale VMware NSX-T, VMware NSX for vSphere, vCenter on VMware Cloud on AWS, VMware SD-WAN by VeloCloud, and Kubernetes deployments.

- Cloud Integration can be done, when VMware Aria Operations for Networks and VMware Aria OperationsVMware Cloud Foundation Operations services are present in the same organization.
- For the cloud integration to be possible, you must have VMware Aria Operations for Networks and VMware Aria OperationsVMware Cloud Foundation Operationsservices in the same geographical location.

This adapter gets problem events from VMware Aria Operations for Networks and publishes the alerts in VMware Aria OperationsVMware Cloud Foundation Operations. Alerts are mapped correctly to the common objects between VMware Aria Operations for Networks and VMware Aria OperationsVMware Cloud Foundation Operations. Common objects supported in this adapter are vCenter Server, VMware NSX-T, and VMware NSX for vSphere. For the common objects, VMware Aria Operations supports launch-in-context to VMware Aria Operations for Networks. This allows the user to perform deep network troubleshooting with the VMware Aria Operations for Networks as the context.

The VMware Aria Operations for Networks adapter only supports VMware Aria Operations for Networks versions 5.2 and above. The VMware Aria Operations for Networks adapter can be installed and configured with the on-prem and SaaS versions of VMware Aria Operations. The VMware Aria Operations for Networks adapter does not support cross platform configuration, it should be on-prem VMware Aria Operations to on-prem VMware Aria Operations for Networks and VMware Aria Operations SaaS to VMware Aria Operations for Networks SaaS.

**Configuring VMware Aria Operations for Networks**

Configure an instance of the VMware Aria Operations for Networks in VMware Aria OperationsVMware Cloud Foundation Operations.

As vCenter and NSX-T are native VMware Aria Operations Management Packs, ensure that you have installed the latest NSX for vSphere Management Pack if you have NSX for vSphere data source configured in VMware Aria Operations for Networks.

In VMware Cloud Foundation Operations, configure the vCenter, NSX-T, and NSX for vSphere Management Pack with the appropriate cloud proxy settings.

1. From the left menu, click **Administration**>**Integrations**.

2. From the **Accounts** tab in the **Integrations** page, click **Add Account** and then select the VMware Aria Operations for Networks card. You can also activate the management pack from the **Repository** tab, where you will find the card in the Available Integrations section. After the management pack is activated, click **Add Account**.
3. Configure the adapter instance.

| Option | Description |
|---|---|
| VRNI FQDN/IP | The FQDN or the IP address of VMware Aria Operations for Networks. |
| Credential | Select and add the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the environment of this management pack, click the plus sign.<br>• **Credential Kind**. Select and configure the Credential Type. You can select either the Local, LDAP, or vIDM network insight credentials.<br>**NOTE**<br>This Management pack supports only the Local, LDAP, and vIDM users that are added in the **User Management** settings of VMware Aria Operations for Networks.<br>– Local - Network Insight Credentials. Enter the credential name, user name of the local user configured in VMware Aria Operations for Networks, and password for that user.<br>– LDAP - Network Insight Credentials. Enter the credential name, LDAP domain configured in VMware Aria Operations for Networks, LDAP user name, and LDAP password for that LDAP user.<br>– vIDM - Network Insight Credentials. Enter the credential name, vIDM FQDN/IP integrated with VMware Aria Operations for Networks, vIDM user name, and vIDM password for that vIDM user.<br>**Credential Name**. Credential Name. |
| Collector / Group | Select the required collector group. |
| Validate Connection | Test Connection should be successful. |

| Option | Description |
|---|---|
| VRNI FQDN/IP | The FQDN or the IP address of VMware Aria Operations for Networks. |
| Credential | Select and add the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the environment of this management pack, click the plus sign.<br>• **Credential Name**. Credential Name. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • **CSP Refresh Token**. The CSP Refresh token of VMware Aria Operations for Networks. To generate the CSP Refresh Token in VMware Aria Operations for Networks:<br>1. Log in to the VMware Cloud Services and select the VMware Aria Operations for Networks account.<br>2. Select your user profile in the top-right corner, and click **My Account**.<br>3. In the **My Account** page, click **API Tokens**, and then click **Generate Token**.<br>4. Click **Generate**.<br>5. Copy or save this token. |
| Collector / Group | Select the default collector group.<br><br>Note: With VMware Aria Operations for Networks, VMware Aria Operations for Networks Management Pack user will not be allowed to select any proxy collector. |
| Validate Connection | To initiate the authentication request to the CSP authentication service, click **Validate Connection**. |

4. The VMware Aria Operations for Networks instance collects events based on common data sources between VMware Aria OperationsVMware Cloud Foundation Operations and VMware Aria Operations for Networks. When you deactivate the **Import problem events as based on common data sources** option, all the events are imported into the VMware Aria OperationsVMware Cloud Foundation Operations.
5. You can collect user-defined events of VMware Aria Operations for Networks as notifications in VMware Aria OperationsVMware Cloud Foundation Operations. To do so, activate the **Import User defined events as Notifications**.
6. Select the severity of the problem events you want to import. By default, all the problem events with moderate and critical severities are imported.
7. Click **Add**.

   The VMware Aria Operations for Networks instance is added to the list.

## NSX

The NSX adapter allows you to retrieve alerts and findings from NSX to VMware Aria OperationsVMware Cloud Foundation Operations.

The NSX adapter supports adapter configuration using vIDM for NSX versions 3.0 and above. The roles and permissions associated with the vIDM users collecting the NSX adapter data is:

| Roles | Permissions |
|---|---|
| Enterprise Admin | Collect all data. |
| Network engineer | • Collect all the NSX resources except the Load Balancer and collect limited routers data. |

*Table continued on next page*

*Continued from previous page*

| Roles | Permissions |
|---|---|
| | Router data collected:<br>– Tier 0 router connected to logical switch.<br>– Tier 1 router created from vCloud Director. |
| • Security Engineer<br>• Security Operator<br>• Auditor | Collect all data except the load balancer. |
| • LB Admin<br>• LB Auditor<br>• Netxpartner Admin | Cannot collect any data. |

**Configuring the NSX Adapter**

To view the roles and permissions associated with the VIDM users collecting the NSX adapter, see NSX Introduction
- To use Principal Identities authentication in VMware Aria OperationsVMware Cloud Foundation Operations you must have created a Principal Identity user in NSX.
- Ensure that you have the client certificate and the key to authenticate the Principal Identity users in VMware Aria OperationsVMware Cloud Foundation Operations.

1. From the left menu, click **Administration › Integrations**.
2. On the **Accounts** tab, click **Add**, and on the **Accounts Types** page, select **NSX**.
3. Configure the adapter instance.

| Option | Action |
|---|---|
| Name | Enter the name for the NSX instance as you want it to appear in VMware Cloud Foundation Operations. |
| Description | Enter any additional information that helps you manage your instances. |

4. Configure the connection.

| Option | Action |
|---|---|
| Virtual IP/ NSX Manager | Enter the FQDN, the IP address, or the Virtual IP of the NSX manager. Both IPv4 address and IPv6 address formats are supported. |
| Credential | Select the credential you want to use to sign on to the environment from the drop-down menu. To add new credentials to access the NSX environment, click the plus sign.<br>1. From the **Manage Credential** dialog, you can click the **Credential Kind** drop-down and select **NSX Client Certificate Crentials** or **NSX Credentials**. The NSX Client Certificate allows you to use Principal Identities user or certificate-based client for |

*Table continued on next page*

*Continued from previous page*

| Option | Action |
|---|---|
| | authentication and the NSX Credential allows you to use local administrator or VMware Identity Manager for authentication.<br>2. Enter the credential details based on your selection. For NSX Credentials, enter the following details.<br>• Credential Name - The name by which you are identifying the configured credentials.<br>• User Name - The user name of the NSX instance.<br>• Password - The password of the NSX instance.<br><br>For NSX Client Certificate Credentials, enter the following details.<br>• Credential Name - The name by which you are identifying the configured credentials.<br>• Client Certificate Data - Enter the value of the client certificate data associated with the principal user.<br>• Client Key Data - Enter the value of the client key data associated with the principal user. |
| Collector /Group | Determine which VMware Cloud Foundation Operations collector or collector group is used to manage the account. If you have multiple collectors or collector groups in your environment, and you want to distribute the workload to optimize performance, select the collector or collector group to manage the adapter processes for this instance.<br>**Review and Accept Untrusted Certificates**<br><br>When you upgrade VMware Cloud Foundation Operations from an earlier version to VMware Cloud Foundation Operations 8.6, the NSX adapter moves to a warning state and the data collection stops. This happens only when you have an adapter that presents a self-signed certificate, or a certificate signed by an untrusted Certification Authority.<br><br>To continue the adapter configuration, you must Validate the Connection, where you are prompted to review and accept the certificate.<br><br>**NOTE**<br>If you have a multi-node cluster configuration, then you are prompted to review and accept the certificate for each node. |
| Advanced Settings | |
| Exclude Resources from Monitoring | Select specific NSX resources that you want to exclude from monitoring in VMware Cloud Foundation |

*Table continued on next page*

*Continued from previous page*

| Option | Action |
|---|---|
| | Operations. Once the exclusion is applied, the monitoring for the specified resources will be disabled.<br><br>**NOTE**<br>Earlier the Management Pack for NSX would collect all the data and monitor all the resource kinds associated with the Management Pack, but now it allows you to exclude specific resources from monitoring. |
| Exclude Services from Monitoring | Select specific NSX management services that you want to exclude from monitoring in VMware Cloud Foundation Operations. Once the exclusion is applied, the monitoring for the specified services will be disabled.<br><br>**NOTE**<br>Earlier the Management Pack for NSX would collect all the data and monitor all the services associated with the Management Pack, but now it allows you to exclude specific services from monitoring. |
| Auto Discovery | Set the **Auto Discovery** to True or False.<br>• **Auto Discovery - True** - Activates auto-discovery of the new objects added to the monitored system. By default, this field is set to True.<br>• **Auto Discovery - False** - Deactivates auto discovery, you must manually discover the objects from the system which you want to monitor . |
| Import Alerts from NSX | You can directly import alerts from VMware NSX (formerly known as VMware NSX-T) into VMware Cloud Foundation Operations by setting the **Import Alerts from NSX** field to **True**.<br>Depending on the number of alerts and the system's performance, importing alerts into VMware Cloud Foundation Operations may take some time.<br><br>By default, this field is set to **False**.<br><br>**NOTE**<br>Direct Alert integration is not supported on VMware Cloud on AWS. |

5. Click **Add**.
   The adapter instance is added to the list.

Verify that the adapter is configured and is collecting data.

**Support for Principal Identities Authentication for the NSX Management Pack**

VMware Aria OperationsVMware Cloud Foundation Operations supports authentication of Principal Identities (PI) using the NSX Management Pack. The Principal Identities (PI) are unique users in NSX who can create an object and ensure that the object can only be modified or deleted by the same identity. The authentication of principal identities is only supported through client certificate. The principal identities authentication is local to NSX Manager, so it does not require VMware Identity Manager, and it is possible to assign a predefined Role-based access control (RBAC) role to the principal identity.

Principal Identities are generally used by third-party applications or cloud management platforms such as Open stack, and Pivotal Container Services (PKS) to ensure that an administrator does not modify the NSX configuration which can generate a mismatch between their view of the NSX environment and the actual configuration.

## Configuring Cloud Federation Adapter

You can configure the adapter instances and add the VMware Cloud Foundation Operations and VMware Cloud Foundation Operations SaaS instances that you want to monitor.

1. From the left menu, click **Administration › Integrations**.
2. Under the Accounts tab, click **Add**.
3. On the Account Types page, click the **Cloud Federation Adapter**.
4. Enter the display name and the description of the adapter.
5. In the **Organization ID** field, enter the VMware Cloud Services ID, if you are creating an instance that monitors VMware Cloud Foundation Operations SaaS.

   **NOTE**
   The Organization ID must be left empty if you are creating an instance that monitors VMware Cloud Foundation Operations.

6. Enter the Host name or the IP address of the VMware Cloud Foundation Operations or VMware Cloud Foundation Operations SaaS source.

   **NOTE**
   While entering the URL for the VMware Cloud Foundation Operations SaaS source, ensure you enter the URL in the format `http://www.host-name.com/ariaops`.

7. Select and add the credential you want to use from the drop-down menu. To add new credentials, click the plus sign.

   From the Credential Kind drop-down:
   - If you want to monitor the VMware Cloud Foundation Operations instances, Select the API Token Credentials and enter the credentials.
   - If you want to monitor VMware Cloud Foundation Operations instances, Select the Principal Credentials and enter the credentials.

8. From the **Collector / Group** drop-down, select Default Collector Group to monitor the VMware Cloud Foundation Operations instances.

   To monitor VMware Cloud Foundation Operations instances, select the appropriate cloud proxies from the **collector / group** drop-down.

9. Click **Validate Connection** to verify if the configuration is successful.
10. Click **Add**.

The Cloud Federation Adapter is added to the list.

## Performing VMware Cloud Director (VCD) Based Multitenancy Operations in VMware Aria OperationsVMware Cloud Foundation Operations

Performing VMware Cloud Director (VCD) Based Multitenancy Operations in VMware Aria Operations
VMware supports its partners to host and sell cloud services built on VMware technology using VMware Cloud Director (VCD), also known as Cloud Director Service (CDS) in the SaaS environment. Cloud Director ServiceVMware Cloud Director provides constructs to segment the virtual infrastructure and offers it as a service to tenants of these partners.

There are several variants of infrastructure that are sold by these partners such as, 'Pay as you go', 'Raw capacity' also known as 'Allocation based', 'Raw capacity with minimum guarantee' also known as 'Reservation based'. The combination of these can be offered to same tenants and it becomes challenging to track usage over a period and charge appropriately. It becomes critical for the cloud service providers as the tenants demand transparency in billing, and the cloud service providers must offer it. Chargeback addresses this by accurately metering the infrastructure. It provides options to configure different models for pricing this metered infrastructure.

> **NOTE**
> The Chargeback capabilities are available only when you have VMware Cloud DirectorCloud Director Service configured in VMware Aria OperationsVMware Cloud Foundation Operations.

This chapter provides information about configuring Chargeback for VMware Cloud DirectorCloud Director Service.

> **NOTE**
> VMware Cloud Director (VCD) is known as Cloud Director Service (CDS) in the SaaS environment.

## Activating Chargeback Capabilities in VMware Aria OperationsVMware Cloud Foundation Operations

Activating VMware Chargeback Capabilities
To activate Chargeback capabilities in VMware Aria OperationsVMware Cloud Foundation Operations, the cloud service providers must install VMware Cloud Director Management Pack.

Perform the following steps to activate the capabilities related to Chargeback in VMware Aria OperationsVMware Cloud Foundation Operations.

> **NOTE**
> If your VMware Aria OperationsVMware Cloud Foundation Operations is integrated with VMware Chargeback in this version of VMware Aria OperationsVMware Cloud Foundation Operations, to activate VMware Aria OperationsVMware Cloud Foundation Operations Chargeback capabilities, you must migrate your data from VMware Chargeback. For more details, see Migration of Chargeback Data to VMware Aria OperationsVMware Cloud Foundation Operations in Case of an Upgrade . After migration, VMware Aria OperationsVMware Cloud Foundation OperationsChargeback capabilities will be activated.

1. From the left menu, click **Administration › Integrations**. Click the **Repository** tab. The page displays installed integrations as clickable cards.
2. Get the VMware Cloud Director adapter from the The Repository Tab.
3. Additionally, you can configure a VMware Cloud Director instance. For more information see, Add an Adapter Instance. Optionally, you can add NSX and VMware Aria Operations Management Pack for Cloud Director Availability adapters. For more information, see Configuring the NSX Adapter.
4. To register the VMware Aria OperationsVMware Cloud Foundation Operations plugin, see Register VMware Aria OperationsVMware Cloud Foundation Operations in a VMware Cloud Director Instance.

All chargeback capabilities are now visible in VMware Aria OperationsVMware Cloud Foundation Operations.

For more details on setting up VMware Aria Operations, see Getting Started with VMware Aria Operations.

## Register VMware Aria OperationsVMware Cloud Foundation Operations in a VMware Cloud Director Instance

You can register VMware Aria OperationsVMware Cloud Foundation Operations in a VMware Cloud Director instance, for tenant view capabilities.

1. Click **Manage Registration** and enter the credentials to register your VMware Aria OperationsVMware Cloud Foundation Operations in the VMware Cloud Director instance.
2. You must complete one of the following:
   - Enter a username and password for a user with VMware Aria OperationsVMware Cloud Foundation Operations register/unregister privileges, or
   - Skip the username and password and select the **Use Collection Credentials** checkbox.

## Migration of Chargeback Data to VMware Aria OperationsVMware Cloud Foundation Operations in Case of an Upgrade

When you upgrade to this version and later versions of VMware Aria OperationsVMware Cloud Foundation Operations, and you have VMware Chargeback, you have the option to switch to VMware Aria OperationsVMware Cloud Foundation OperationsChargeback and use all the capabilities from one place. This can be achieved through the migration process. The existing settings and configurations will migrate to VMware Aria OperationsVMware Cloud Foundation Operations.

**Steps to Migrate Chargeback Data to VMware Aria Operations**

1. From the VMware Chargeback UI, in the left menu, click **Migration**.
2. From the **Migration** page, you will see the list of tasks to be migrated with the Start and End time. You can also view the status of each task that is being migrated. Click **Initiate Migration**.

   **NOTE**
   Click on each task row to view information about the number of objects that have been migrated, the number of objects that have been ignored, the migrated resources, error messages if any, and so on.

List of tasks in Chargeback and their corresponding location in VMware Aria OperationsVMware Cloud Foundation Operations

**Table 112: Migration Details from Chargeback to VMware Aria OperationsVMware Cloud Foundation Operations**

| Chargeback Migration Steps and Location in Chargeback | Location in VMware Aria Operations |
|---|---|
| **Initialization** | NA |
| **Email Configuration**<br>**Admin Setting › Configure › Configure Email** | Click **Operations › Configurations**, and then click the **VCD Tenant Email Configuration** tile. |
| **Pricing Policies**<br>**Pricing › Configuration** | Click **Operations › Configurations**, and then click the **Policy Definition** tile.<br>**Pricing Policies Assignment**<br><br>All pricing policies from Chargeback will be converted to separate policies under the default Policy in VMware Aria OperationsVMware Cloud Foundation Operations. They will have the same name as they have in the Chargeback |

*Table continued on next page*

*Continued from previous page*

| Chargeback Migration Steps and Location in Chargeback | Location in VMware Aria Operations |
|---|---|
| | UI with the pricing card UUID added as a suffix. After migration Active Polices are the ones that have at least one object assignment in Chargeback.<br>**Pricing Policies Assignment Migration**<br>• If there is just one active policy in VMware Aria OperationsVMware Cloud Foundation Operations before migration, then all the policy assignments from Chargeback will be migrated to VMware Aria OperationsVMware Cloud Foundation Operations accordingly.<br>• If there is more than one active policy in VMware Aria OperationsVMware Cloud Foundation Operations, then after migration, the policy assignment of Organization and Organization VDC objects should be done manually. In this scenario it is recommended to proceed with policy assignments immediately to ensure continuous pricing calculations. |
| **Bill Schedules** | Bill schedules migration should be done manually. For more details about scheduling bills, see Scheduling Bill Generation Using Automation Central. |
| **Notification Rules**<br>**Alerts › Tenant Alerts › Notification Rules** | Click **Operations › Configurations**, and then click the **VCD Tenant Notifications** tile. |
| **Report Schedules**<br>**Reports › Tenant Reports** | Click **Operations › Reports › Manage**<br>Filter by **VCD Tenants** and then view the **Schedules** column to view the Scheduled reports. |
| **Tenant Reports**<br>**Reports › Tenant Reports** | Click **Operations › Reports › Generated Reports**<br>Filter by **VCD Tenants** and view the generated reports. |
| **Generated Bills**<br>**Pricing › Bills** | Click **Cost › Bills** |
| **Bill Retention Settings**<br>**Admin Settings › Settings › Data Retention** | Click **Administration › Control Panel › Chargeback › Data Retention › Generated Bills**. |
| **Finalization**<br>**Admin Settings › Support › Self Health** | Click **Administration › Control Panel › Audit › Chargeback Migration Audit**<br>• Users created with the role of VCD Tenant Admin will be removed.<br>• All VCD instances which had VMware Chargeback registered will now be switched to VMware Aria OperationsVMware Cloud Foundation Operations.<br>– Manual registration may be required in certain cases. These instances will be reported in the Finalization step summary.<br>• VMware Aria OperationsVMware Cloud Foundation OperationsChargeback capabilities will be activated.<br>• A detailed report of migration can be found at **Administration › Control Panel › Audit › Chargeback Migration Audit**. For more details, see Chargeback Migration Audit. |

# Managing Chargeback Administration Settings

Managing Chargeback Administration Settings
Cloud service providers can manage the access of metrics and pages by configuring specific metrics or pages that the tenants are permitted to access. Cloud service providers can also define storage pricing policies and manage historical data related to data retention in VMware Aria OperationsVMware Cloud Foundation Operations.

Ensure that you have activated VMware Cloud Director, for details, see Activating Chargeback Capabilities in VMware Aria Operations.

### Managing Access to Metrics

Cloud service providers can manage the access of metrics to the tenants in VMware Aria OperationsVMware Cloud Foundation Operations. You can configure specific metrics that your tenants are permitted to access.

1. To manage access of metrics, from the left menu, click **Administration › Control Panel**, and then click the **Chargeback** tile.
2. In the **Manage Metrics** tab, select an object type for which you want to activate or deactivate metric access from the **Select Object Type** list.

   **NOTE**
   By default, all metrics are deactivated.

3. To activate or deactivate a particular metric, select the metric and select **Activate** or **Deactivate** from the **Actions** drop-down list. You can also activate or deactivate metric by clicking the drop-down arrow against a specific metric and selecting **Activate** or **Deactivate**.

   **NOTE**
   To activate or deactivate all metrics, click the Select All icon and select **Activate** or **Deactivate** from the **Actions** drop-down list.

4. Click **Save**.

### Managing Access to Pages

Cloud service providers can manage the access of pages to the tenants in VMware Aria OperationsVMware Cloud Foundation Operations. You can configure specific pages that your tenants are permitted to access.

1. To manage access of pages, from the left menu, click **Administration › Control Panel**, and then click the **Chargeback** tile.
2. Click the **Manage Pages** tab, and select the pages for which you want to grant access to the tenant.
3. Click **Save**.

### Defining Storage Pricing

Cloud service providers can charge for storage either based on storage policies or independent of it in VMware Aria OperationsVMware Cloud Foundation Operations.

Storage prices are calculated based on usage by storage policies that are independent of underlying VMs and templates, or by aggregating the usage from underlying VMs and templates. The difference is that the former considers indirect disks such as log disks and swap disks, whereas the latter considers only the storage used directly by the VMs.

1. To define price setting, from the left menu, click **Administration › Control Panel**, and then click the **Chargeback** tile.
2. Click the **Price Settings** tab, and select an option to aggregate storage charges for PAY AS YOU GO (PAYG) org-VDCs.
   a) To charge independent of storage policies, select **Aggregate storage charges from VMs, Media, vAPP templates and independent disks for PAYG org-VDCs**.
   b) To charge based on storage policies, select **Aggregate storage charges from Storage profiles for PAYG org-VDCs**.
3. Click **Save**.

**NOTE**
The **Default Base Rate** option under **Policy Definition › VCD Pricing › Storage Rate** is considered only when you have selected Aggregate storage charges from Storage profiles for payg Org-VDCs in the **Pricing Settings** tab.

### Managing Historical Data

You can view and manage historical data retention related to system settings. You can tune these settings to manage storage space consumed by VMware Aria OperationsVMware Cloud Foundation Operations.

1. To view and manage historical data, from the left menu, click **Administration › Control Panel**, and then click the **Chargeback** tile.
2. Click the **Data Retention** tab, and enter the number of days for which you want to retain the generated bills.

**NOTE**
You can retain the generated bills to a maximum of 1440 days.

3. Click **Save**.

## Configuring VCD Pricing Details

You can edit the parent policy settings and configure your OVDC settings using the **VCD Pricing** tab in VMware Aria OperationsVMware Cloud Foundation Operations.

If you want to copy the VCD pricing settings from the policy currently being edited to another policy, click **Copy local changes to other policy** and select the policy to which you want to copy the settings. The copied pricing configuration will override any existing local pricing configuration in the target policy.

1. From the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile.
2. Select the required policy or click **Add** to add a new policy.
3. Select the required policy, and in the right pane, click **Edit Policy**.
4. In the <policy name> [Edit] workspace, click the **VCD Pricing** card.
5. Click the Lock icon to override parent policy settings.
6. Select if you want to activate or deactivate the pricing engine.
7. Select the pricing policy type from the drop-down menu. The pricing policy type determines your billing model based on the Organization VDC type.
8. Configure the OVDC settings.
   - **Compute Rate**: Click the Lock icon to edit the parent policy settings. Select if you want to charge based on the Sizing Policy or based on CPU/Memory Rate.

The following options appear if you select **CPU/Memory Rate**.

- CPU Count Based/CPU GHz Based: Select if you want to charge the CPU rate based on GHz or CPU count.
- Charge Period: The Charge Period indicates the frequency of charging.
- Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
- Default Base Rate (per CPU count): Enter a valid number for Default Base Rate.
- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.

  **NOTE**
  This field appears when you charge the CPU based on GHz.

- Default Base Rate (per GHz): Enter a valid number for Default Base Rate.

  **NOTE**
  This field appears when you charge the CPU based on GHz.

- Other Base Rate Slabs: Using slabs, you can optionally charge different rates depending on the number of CPUs used. Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 5 in Greater than or equal and 5 as Base Rate, it means if the usage is 5 CPU and above, then the Base Rate of 5 will be applied for the whole usage.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- Fixed Cost (per charge period per VM): Enter a valid number. Fixed costs do not depend on the units of charging.

Enter details for **Memory Rate**:

- Charge Period: The Charge Period indicates the frequency of charging.
- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
- Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
- Default Base Rate (per GB): Enter a valid number for Default Base Rate.
- Base Rate Slab: Using slabs, you can optionally charge different rates depending on the memory allocated. Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 50 in Greater than or equal and 10 as the Base Rate, it means if the usage is 50 GB and above, then the Base Rate of 10 will be applied for the whole usage.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- Charge Overage: This rule applies only to the allocation pool model, where guaranteed resources are some percent of total available resources. In this case, for the guaranteed percent, a normal base rate will be applied. If usage goes beyond the guaranteed percent, the rate mentioned in overage will be applied (for the delta usage which is higher than guaranteed).

- Overage Memory Rate: Enter an overage memory rate. This field appears only when you activate the **Charge Overage** option. If usage goes beyond the guaranteed percent, the rate entered in overage will be applied (for the delta usage which is higher than guaranteed).
- Fixed Cost (per charge period per VM): Enter a valid number. Fixed costs do not depend on the units of charging.

The following options appear if you select **Sizing Policy**. Sizing policies are a way of defining template VM sizes such as Small, Medium, and Large, in terms of CPU and Memory.

- Add: Click **Add** and enter the following details:
  - Sizing Policy Name: Select the sizing policy name from the drop-down.
  - Charge Period: The Charge Period indicates the frequency of charging.
  - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
  - Base Rate: Enter a base rate.

  The slabs that you add appear in the table below. To edit, delete, or add new slabs, click the vertical ellipses and select the desired option.

- **Storage Rate**: Click **Storage Rate** in the left pane and then click the Lock icon to edit the parent policy settings.
  - Click **Create Storage Policy** and enter the following details:
    - Storage Policy Name: Select a storage policy name from the drop-down list.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
    - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
    - Default Base Rate (per GB): Enter a default base rate.

      **NOTE**
      The Default Base Rate option is considered only when you have selected **Aggregate storage charges from Storage profiles for payg Org-VDCs** option under **Administration › Control Panel › Chargeback › Pricing Settings** tab.

    The storage policy that you add appears in the table below. To edit, delete, or add new slabs, click the vertical ellipses and select the desired option. Using slabs, you can optionally charge different rates depending on the storage allocated.

- **Network Rate**: Click **Network Rate** in the left pane and then click the Lock icon to edit the parent policy settings.
  - **Network Data**
    - External Network Transmit (per MB): Enter the rates for external network transmit.
    - External Network Receive (per MB): Enter the rates for external network receive.
  - **Network Transmit Rate (Bandwidth)**
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Usage: Select the usage based on which you want to charge.
    - Default Base Rate (per MBps): Enter a default base rate.
    - Add Slab: Using slabs, you can optionally charge different rates depending on the network data consumed. Click **Add Slab** and enter the following details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.
      - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

      The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 50 in Greater than or equal and 10 as the Base Rate, it means if the usage is 50 Mbps and above, then the Base Rate of 10 will be applied for the whole usage.

The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Network Receive Rate (Bandwidth)**
  - Charge Period: The Charge Period indicates the frequency of charging.
  - Usage: Select the Usage based on which you want to charge.
  - Default Base Rate (per MBps): Enter a default base rate.
  - Add Slab: Using slabs, you can optionally charge different rates depending on the network data consumed. Click **Add Slab** and enter the following details:
    - Greater than or equal: Enter a valid number.
    - Base Rate: Enter a base rate.
    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The base rate slab allows you to change the rate of charging based on the resources used. For example, if you enter 50 in Greater than or equal and 10 as the Base Rate, it means if the usage is 50 Mbps and above, then the Base Rate of 10 will be applied for the whole usage.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **Advanced Network Rate**: Click **Advanced Network Rate** in the left pane and then click the Lock icon to edit the parent policy settings.
  - Edge Gateway Size: Enter Base Rate, Period, and Fixed Cost for sizes Compact, Large, X-Large, and Quad Large.
    Chargeback allows you to define the size of the edge gateway and assign differential prices based on the edge size.
  - Distributed Firewall: Enter Charge Period and Base Rate (per count).
  - Edge Services: Enter Charge Period and Base Rate (per count) for IP Count, and enter Charge Period and Base Rate (if enabled) for HA, DHCP, IPV6, IP Sec, LB, NAT, SSL VPN, L2 VPN, Firewall, Static Routing, BGP Routing, and OSPF Routing.
    IP Count is the unique IP count available on the external network of the Org-VDC. Pricing can be performed based on the count of these IPs.

    Apart from the basic data transfer, there are additional value-added services offered in VMware Cloud DirectorCloud Director Service in combination with NSX. All the network services associated with specific edges such as HA, DHCP, IPV6, IP Sec, Load Balancer, NAT, SSL VPN, L2 VPN, Firewall, Static Routing, BGP Routing, and OSPF Routing are considered for charging based on whether these services are activated or not. If services are activated for a specific day and a base rate is applied for that service, then that particular service gets charged for that specific day. If the service is deactivated on any day then the base rate will not be applied.
  - Network Service Pricing (NSXT Only)
    - Firewall Charges (per firewall rule count): Enter Charge Period and click **Add Slab** to enter the following base rate slab details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.
      - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

      The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
    - L2VPN charges (per L2VPN count): Enter Charge Period and click **Add Slab** to enter the following base rate slab details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.
      - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- Load Balancer Charges (per load balancer count): Enter Charge Period and click **Add Slab** to enter the following base rate slab details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- NSX Advanced Load Balancer (Throughput): Enter Charge Period, Usage, Default Base Rate (per MBps), and click **Add Slab** to enter the following base rate slab details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  Cloud service providers can charge for consumption of NSX Advanced Load Balancer based on throughput. You can create slabs with different rates for throughput values.

  **NOTE**
  The unit of charging for 'throughput' is 'mbitsps'.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

  **NOTE**
  The default base rate will be applied for any range of usage that is not covered in the above set of slabs.

- **Guest OS Rate**: Click **Guest OS Rate** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for any additional charges that are to be applied on Virtual machines based on their Discovered Guest Operating System by vCenter.
  - Click **Create Guest OS Rate** and enter the following details:
    - Guest OS Name: Enter a guest OS name.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
    - Base Rate: Enter a base rate.

    The guest OS rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **Cloud Director Availability**: Click **Cloud Director Availability** in the left pane and then click the Lock icon to edit the parent policy settings.
  - Per Replica Charge: This section is used to set pricing for replications created from Cloud Director Availability. You can charge for each replication object, based on the SLA profile they belong to. For charging replications without any SLA Profile assigned, please enter None as the SLA Profile name. Click **Create Per Replication Charge** and enter the following details:
    - Replication SLA Profile Name: Enter a profile name.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Base Rate: Enter a base rate.

    The replication charges that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

  - Storage Usage Charge: This section is used to set additional pricing for storage used by Cloud Director Availability replications in Cloud Director. The storage usage defined in this tab will be added additionally to the Storage Policy Base Rate.
    Click **Create Storage Usage Charge** and enter the following details:
    - Storage Policy Name: Select a storage policy name from the drop-down list.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Base Rate: Enter a base rate.

    The storage charges that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **vCenter Tag Rate**: Click **vCenter Tag Rate** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for any additional charges that have to be applied to the VMs based on their discovered Tags from vCenter.
  - **Fixed Rate**: Click **Create Base Rate** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
    - Base Rate: Enter a base rate.

    The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **Alternate Pricing Policy**: Click **Create Base Rate** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Pricing Policy: Select the alternate pricing policy name.
    - Priority: Select the priority for the alternate pricing policy. When the metadata or tag-based charges overlap, setting a priority allows you to define which policy should be processed first.

    The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **VCD Metadata Rate**: Click **VCD Metadata Rate** in the left pane and then click the Lock icon to edit the parent policy settings. The metadata based prices are available in bills only if the Enable Metadata option is enabled in the VMware Aria Operations Management Pack for VMware Cloud Director. This section is used to charge for any additional charges that have to be applied to the VMs based on their discovered metadata from the Cloud Director.
  - **Fixed Rate**: Click **Create Base Rate** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based on Power State: This decides whether the charge should be applied based on the power state of the VM.
    - Base Rate: Enter a base rate.

    The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **Alternate Pricing Policy**: Click **Create Base Rate** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Pricing Policy: Select the alternate pricing policy name.
    - Priority: Select the priority for the alternate pricing policy. When the metadata or tag-based charges overlap, setting a priority allows you to define which policy should be processed first.

    The base rates that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

- **One Time Fixed Cost**: Click **One Time Fixed Cost** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for one time incidental charges on Virtual machines, such as creation/setup charges, or charges for one-off incidents like installation of a patch. These costs do not repeat on a recurring basis.
  - VM Creation: Enter the one time fixed cost.
  - VCD Metadata: Click **Create VCD Metadata** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - One Time Fixed Cost: Enter the one time fixed cost.

    The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - vCenter Tag: Click **Create vCenter Tag** and enter the following details:
    - vCenter Tag Key: Enter a tag key.
    - vCenter Tag Value: Enter a tag value.
    - One Time Fixed Cost: Enter the one time fixed cost.

    The details that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **Rate Factors**: Click **Rate Factors** in the left pane and then click the Lock icon to edit the parent policy settings. Use Rate Factors to either increase or discount the prices against individual resources consumed by the Virtual Machines or by whole charges against the Virtual Machine.
  - VCD Metadata: Click **Create VCD Metadata** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Change the price of: Select the required option from the drop-down list.
    - By applying a factor of: Enter a valid number. For example, if you want to increase the price of the CPU which has a tag 'Tag1-Value1' by 20% then select CPU from the **Change the price of** drop-down menu and enter 1.2 in **By applying a factor of**.

    The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - vCenter Tag: Click **Create vCenter Tag** and enter the following details:
    - vCenter Tag Key: Enter a tag key.
    - vCenter Tag Value: Enter a tag value.
    - Change the price of: Select the required option from the drop-down list.
    - By applying a factor of: Enter a valid number. For example, if you want to increase the price of the CPU which has a tag 'Tag1-Value1' by 20% then select CPU from the **Change the price of** drop-down menu and enter 1.2 in **By applying a factor of**.

    The details that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **Tanzu Kubernetes Clusters**: Click **Tanzu Kubernetes Clusters** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for Tanzu K8s clusters and objects below an Org VDC based on certain attributes of K8s like CPU, Storage, Memory, and so on.
  - **Cluster Fixed Cost**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Fixed Cost (per count): Enter a valid number.
  - **Cluster CPU Rate**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
    - Default Base Rate (per GHz): Enter a default base rate.
    - Base Rate Slab: Click **Add Slab** and enter the following details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.

    - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

    The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **Cluster Memory Rate**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
    - Default Base Rate (per GB): Enter a default base rate.
    - Base Rate Slab: Click **Add Slab** and enter the following details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.
      - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

      The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **Cluster Storage Rate**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Default Base Rate (per GB): Enter a default base rate.
    - Base Rate Slab: Click **Add Slab** and enter the following details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.
      - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

      The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

      **NOTE**
      The default base rate will be applied to any range of usage that is not covered in the above set of slabs.

- **CSE Kubernetes Cluster**: Click **CSE Kubernetes Cluster** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to charge for K8s clusters and objects below an Org VDC based on certain attributes of K8s like CPU, Storage, Memory, and so on.
  - **Cluster Fixed Cost**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Fixed Cost (per count): Enter a valid number.
  - **Cluster CPU Rate**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
    - Default Base Rate (per GHz): Enter a default base rate.
    - Base Rate Slab: Click **Add Slab** and enter the following details:
      - Greater than or equal: Enter a valid number.
      - Base Rate: Enter a base rate.
      - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

      The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  - **Cluster Memory Rate**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.

- Default Base Rate (per GB): Enter a default base rate.
- Base Rate Slab: Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

– **Cluster Storage Rate**: Enter the following details:
- Charge Period: The Charge Period indicates the frequency of charging.
- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
- Default Base Rate (per GB): Enter a default base rate.
- Base Rate Slab: Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

  The slabs that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

  **NOTE**
  The default base rate will be applied to any range of usage that is not covered in the above set of slabs.

- **Additional Fixed Cost**: Click **Additional Fixed Cost** in the left pane and then click the Lock icon to edit the parent policy settings. This section is used to add any other costs that have to be applied at the Org-VDC level. This can be used for charges such as overall tax, overall discounts, and so on. These can be applied to specific Org-VDCs based on Org-VDC metadata.
  – **Fixed Cost (is applied at Org-VDC level)**: Enter the following details:
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Fixed Cost: Enter a valid number.
  – **VCD Metadata (Addition cost is applied at Org-VDC level)**: Click **Create VCD Metadata** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - Charge Period: The Charge Period indicates the frequency of charging.
    - Price: Enter a valid price.

    The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
  – **VCD Metadata One Time (Addition cost is applied at Org-VDC level)**: Click **Create VCD Metadata One Time** and enter the following details:
    - Metadata Tag Key: Enter a tag key.
    - Metadata Tag Value: Enter a tag value.
    - One Time Fixed Cost: Enter a one time fixed cost.

    The metadata that you add appears in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.
- **Organization**: Click **Cloudian Storage** in the left pane and then click the Lock icon to edit the parent policy settings. A third-party storage service that allows the cloud service providers to provide additional storage to the tenants and charge them for the same. You can charge for the consumption of cloudian storage by creating slabs with different rates for storage values.
  Enter the following details:
  - Charge Period: The Charge Period indicates the frequency of charging.

- Charge Based On: The Charge Based On indicates the pricing model based on which the charge is applied.
- Default Base Rate (per GB): Enter a default base rate.
- Base Rate: Click **Add Slab** and enter the following details:
  - Greater than or equal: Enter a valid number.
  - Base Rate: Enter a base rate.
  - Explanation: Displays an explanation of the slab based on the values entered in the above fields.

The details that you add appear in the table below. To edit or delete the entries, click the vertical ellipses and select the desired option.

**NOTE**
The bills for cloudian storage will be generated only at the organization level.

9. Click **Preview** to review your changes and then, click **Save**.

You can assign policies to the required vCenter/Cluster Compute Resource under **Policy Assignment**. Navigate to **Operations › Configurations**, and then click the **Policy Assignment** tile. For details, see Assigning Policies.

**NOTE**
For details on configuring vCenter pricing, see Configuring vCenter Pricing Details.

## Creating Notification Rules for Tenant in VMware Aria OperationsVMware Cloud Foundation Operations

Creating Notification Rules for Tenant
Cloud service providers can add, manage, and edit your notification rules for tenants in VMware Aria OperationsVMware Cloud Foundation Operations.

Before you can create and manage your notification rules, you must configure the outbound alert plug-in instances. For details on configuring outbound plug-ins, see Adding Outbound Notification Plug-Ins.

**NOTE**
Chargeback allows you to only add a Standard Email Plug-In for outbound alerts.

1. To manage your notifications, from the left menu, click **Operations › Configurations**, and then click the **VCD Tenant Notifications** tile. On the toolbar, click **Add** to add a rule, or click the vertical ellipsis and select **Edit** to edit the selected rule.
2. Enter the following notification details.

| Option | Description |
|---|---|
| Name | Provide a name for the notification rule. |
| Description | Describe the notification rule. |
| Notification Status | Either activate or deactivate a notification setting. Deactivating a notification will stop the alert notification for that setting and activating it will activate the alert notification. |

3. Click **Next**.

4. Select if you want to send notification at the organization level or at a specific OVDC level. The list of OVDCs/ Organizations is updated based on your selection.
5. From the list of OVDCs/Organizations, drag and drop the OVDC or Organization that should receive a notification, into the left pane.

   **NOTE**
   You can only add those OVDCs that have email addresses and email outbound instances configured. To configure email addresses and email outbound instances, see Configuring Tenant Email and Add a Standard Email Plug-In for Outbound Alerts respectively.

   The notifications will be sent to all selected OVDCs with configured email addresses and email outbound instances.
6. Click **Next**.
7. Define criteria for the notification rule. From the list of alert definitions, drag and drop the alert definition for which you want to trigger a notification, into the left pane.

   **NOTE**
   By default, the filter Only Tenant is set to **Yes**.

   The notification will be sent when ANY of the selected alert definitions triggers an alert.
8. Click **Create**.

You can view the created notification from **Operations › Configurations**, and then click the **VCD Tenant Notifications** tile.

## Configuring Tenant Email in VMware Aria OperationsVMware Cloud Foundation Operations

Configuring Tenant Email
After creating an email outbound, cloud service providers can assign email addresses to Org-VDCs. These email addresses receive alerts and reports related to a particular Org-VDC.

Ensure that you have created an email outbound. For details, see Add a Standard Email Plug-In for Outbound Alerts.

1. To configure tenant email, from the left menu, click **Operations ›  Configurations**, and then click the **VCD Tenant Email Configuration** tile.
2. Select if you want to configure tenant email at the organization level or at a specific OVDC level. The list of OVDCs/ Organizations is updated based on your selection.
3. Select the required OVDCs/Organization, and on the toolbar, click **Configure** to add or edit email configuration.
4. Enter the email address of tenant to whom the email has to be sent.
5. Select an email outbound to which you want to send the email.
6. Click **Save**.

## Managing Chargeback Reports

Cloud service providers can view, schedule, and generate reports in VMware Aria OperationsVMware Cloud Foundation Operations. You can create report templates for your tenants.

A report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your environment.

As a cloud service provider, you can:
- View and access all report templates, for information on accessing reports, see Accessing Report Templates .

- Create report templates for tenants by selecting the **Template for Tenant** option, for more information, see Create a Report Template.
- Generate a report for a specific Org-VDC, for more information, see Generate and Regenerate a Report. You can order the reports by the date and time that they were created, the report name, the object on which the report is generated, the owner, or the status.
- Schedule a report for a specific Org-VDC, for more information, see Schedule a Report. Cloud service providers can generate tenant reports on demand or they can schedule them to be emailed to their tenants.

## Generating Bills in VMware Aria OperationsVMware Cloud Foundation Operations

Generating Bills
Chargeback allows you to generate monthly bills to provide an account of the overall expenses used for resources in an organization.

As a cloud service provider, you can charge different OVDCs and generate bills at the Organization/OVDC level. You can provide a bill summary at organization level by separating the costs incurred by the different OVDCs. You can also charge for resources consumed at the organization level rather than the OVDC level.

> **NOTE**
> There is a service in Chargeback for automatic bill generation. On the first of every month, bills are generated for all the available organization VDCs for the previous month (1st to 30th or 31st). The generated bills are listed under **Cost › Bills**.

1. To generate a bill, from the left menu, click **Cost › Bills**, and then click **Generate Bill**.
2. Enter the following details.

| Option | Description |
|---|---|
| Name | Enter a name for the bill. |
| Description | Describe the bill. |
| Start Date | Select a start date for the bill. |
| End Date | Select an end date for the bill. |
| Advanced Settings | |
| Policy | If you want to generate a bill using a specific policy, then select a policy from the list. Only the policies that have VCD pricing configured will appear in the list.<br>It is not mandatory to select a policy from the list. If you do not select any policy, then the bills will be calculated based on VCD pricing configurations present in the default policy. |

3. Click **Next** to select resources.
4. Select if you want to generate a bill at the organization level or at a specific OVDC level. The list of resources is updated based on your selection.

> **NOTE**
> If you select **Organization**, then all the OVDCs under it are selected automatically. Deselecting any of the OVDCs will undo the selection at the Organization level.

5. From the list of resources, drag and drop the OVDC or Organization into the left pane.

**NOTE**
You can use filters to filter resources by OVDC Name and Organization.

6. Click **Create**.

You can view the bill that you generated from **Chargeback › Bills**.

For bills generated at the Organization level, you can view the bill summary of the Organization and associated Organization VDCs. The Cloudian Storage section is available only when cloudian storage is configured and applied as part of the pricing policy.

## Scheduling Bill Generation Using Automation Central

You can create a schedule for bill generation using the Automation Central feature of VMware Aria OperationsVMware Cloud Foundation Operations.

Make sure you have VMware Cloud Director configured in VMware Aria OperationsVMware Cloud Foundation Operations to configure the pricing policy for the bill generation. For more information, see the topic, Performing VMware Cloud Director (VCD) Based Multitenancy Operations in VMware Aria OperationsVMware Cloud Foundation Operations

1. Click **Operations › Automation Central** in the left menu of VMware Aria OperationsVMware Cloud Foundation Operations. Alternatively, click **Cost › Bills › Schedule Bills** from the three horizontal dots next to the **GENERATE BILL** button on the **Bills** page.
2. In the Automation Central page, click **Add Job**.
   The **Create New Job** page opens. This page displays the following cards:
   - Reclaim
   - Rightsize
   - Billing. This card is available when you when you have VMware Cloud Director configured in VMware Aria OperationsVMware Cloud Foundation Operations.
   - Actions
3. Click the **Billing** card.
4. In the **Billing Information** step of the wizard, provide the following inputs and then click **Next**:

| Property | Description |
|---|---|
| Name | Specify a name for the job. This is displayed in the calendar. |
| Description | Provide a description for the job. |

5. In the **Select an Object** step of the wizard, use the toggle button to display either OVDC or Organization objects in the right hand pane. Drag the OVDC or Organization objects to the left hand pane to generate bills for those objects. You can select multiple objects. Click **Next**.
6. In the **Schedule** step of the wizard, provide the following inputs:

| Property | Description |
|---|---|
| Start Date | Set a start date which is on, or after the current date, for the job to start. Select a date from the calendar displayed on the page. |

*Table continued on next page*

*Continued from previous page*

| Property | Description |
|---|---|
| | **NOTE**<br>There will be a delay of up to five minutes for a job to run. The actual start time for an action also depends on the action itself and the number of target objects involved. |
| Time zone | For the start time and date to be calculated for the job, select a time zone. Choices are:<br>• Browser - based on the current location reported by the browser.<br>• Host - based on the current location of the host machine.<br>• GMT based time zone - Based on a location calculated as per UTC+0. |
| Start Time | Select a start time for the job. The drop-down menu provides options in five minute intervals, starting closest to your current browser reported time. |
| Recurrence | Set a recurrence for the job. The choices are:<br>• **Weekly**. Select the days of the week when you want the job to run, by clicking the day of the week abbreviation.<br>• **Monthly**. Select the months when you want the job to run, by clicking the month abbreviation. By default, all the months are selected. For the months that you select, you can configure the job to run:<br>– On specific days of the month, by clicking the number. Or, the last day of the month, without specifying the exact date, by clicking **Last**.<br>– On the specific number of the week (first, second, third, fourth or the last) in the month, combined with the specific day of the week.<br>You can make multiple selections for each of the options in the drop down menu.<br>For the **Daily**, **Weekly**, **Monthly** options, you can set when the job must end based on a date, or the number of occurrences. |
| End | Provide a date when the schedule will end. |
| Bill Start Day | From the drop down options, provide the start date for the billing period. |
| Bill End Day | From the drop down options, provide the end date for the billing period. |
| Notifications | In the **Notifications** section, select the **Receive Updates on Job via Email** check box to receive notifications two hours before the job is set to run and after it has been executed. For the email to be sent, you must also select the email outbound plugin from the drop down menu, and enter the email address to which |

*Table continued on next page*

*Continued from previous page*

| Property | Description |
|---|---|
| | the email must be sent. If you have not created an email outbound plugin, see the topic, Outbound Settings in .<br>**NOTE**<br>You can send updates via email to only one email ID. |

7. Click **Create** to complete the steps in the wizard and create the job.

The job to create a bill as per the schedule is created and added to the Automation Central calendar. When you click **Automation Central** in the left menu, you see the calendar of events. You can also switch the tabs to **History**, **Jobs** to see more information about the jobs.

**NOTE**
The **Reports** tab does not display any information related to billing.

## Viewing Chargeback Summary

The **Chargeback Overview** page provides an overview of resources of an organization in VMware Aria OperationsVMware Cloud Foundation Operations.

From the left menu, click **Cost › Chargeback** to view the **Chargeback Overview** page.

| Widget | Description |
|---|---|
| Summary | • Provides the number of organization VDCs, vApps, virtual machines, and number of running virtual machines that are available in an organization.<br>• Provides total cost of the entire data center for a given OVDC, and ongoing price in US$. |
| Capacity Overview | Provides an overview of CPU, memory, and storage use of an organization. |
| Organization Details | Provides detailed information about each organization, such as the number of organization VDCs, vApps, VMs, running VMs, cost in US$ charged to the cloud service provider, and ongoing price in US$ charged to the tenant.<br>**NOTE**<br>Click **Export** to export organization details as an Excel file. |

## Optimizing Capacity and Improving Performance in VMware Cloud Foundation Operations

Optimizing Capacity and Improving Performance
In VMware Cloud Foundation Operations, Capacity Optimization is a key feature that focus on efficiently utilizing and managing resources within a virtualized infrastructure, such as a VMware vSphere environment. These features help organizations maintain optimal performance and maximize the utilization of their resources while ensuring that capacity remains within acceptable limits.

## Capacity Optimization

This aspect of VMware Cloud Foundation Operations involves forecasting and planning for future resource needs based on historical data and current usage patterns. The goal is to ensure that the virtual environment has the necessary resources to accommodate workload growth without sacrificing performance. Capacity optimization involves activities like identifying underutilized or overprovisioned resources, predicting when additional resources will be needed, and suggesting right-sizing of virtual machines and hosts.

## Performance Improvement

Performance improvement focuses on enhancing the efficiency and responsiveness of the virtual environment. It involves monitoring the performance metrics of virtual machines, hosts, storage, and other components to detect performance bottlenecks, latency issues, and other concerns that could impact the overall performance of applications running on the virtualized infrastructure. .

## Capacity Optimization Concepts

Capacity optimization in VMware Cloud Foundation Operations involves various key concepts and practices to ensure optimal resource utilization, plan for growth, and maintain performance within a virtualized environment.

### How Does Capacity Optimization Work in VMware Aria Operations

How Does Capacity Optimization Work
Capacity Optimization in VMware Aria OperationsVMware Cloud Foundation Operations is achieved using powerful integrated functions - capacity overview, workload balancing and optimization, repurposing of underutilized resources, and what-if predictive scenarios - to reach optimal system performance.

Capacity planners must assess whether physical capacity is sufficient to meet current or forecasted demand. With robust capacity planning and optimization, you can manage your production capacity effectively as your organization addresses changing requirements. The objective of strategic capacity optimization is to reach an optimal level where production capabilities meet ongoing demand.

VMware Aria OperationsVMware Cloud Foundation Operations analytics provide precise tracking, measuring and forecasting of data center capacity, usage, and trends to help manage and optimize resource use, system tuning, and cost recovery. The system monitors stress thresholds and alerts you before potential issues can affect performance. Multiple pre-set reports are available. You can plan capacity based on historical usage, and run what-if scenarios as your requirements expand.

#### Capacity Optimization Using Overview, Reclaim, Workload Optimization and What-If

The Capacity Optimization provides four integrated functions - Overview, Reclaim, Workload Optimization, and What-If Scenarios - that give an overview of the status of all data center activity and trending. You can conduct on-the-spot analysis, including drilling down into further detail on any object to identity possible performance problems or anomalies. You can rebalance and optimize compute resources. The system further identifies underutilized workloads (virtual machines) and calculates the potential cost savings that can accrue when these resources are reclaimed to be deployed more effectively. You can interact with and manipulate data and outcomes based on your requirements.

Use the Capacity Optimization and Reclaim features to assess workload status and resource contention in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

Workload Optimization provides for moving virtual workloads and their file systems dynamically across datastore clusters within a data center or custom data center. You can potentially automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention triggers an alert and automatically runs an action, a data center performs at optimum.

In addition, the What-If Analysis function- can run scenarios that help determine where additional system resources can be brought online.

> **NOTE**
> You may see a data center or cluster labeled as optimized when it has few or no days remaining before CPU, memory, or storage is predicted to run out. That is because these are two different measures of data center and cluster health. A data center can be running at optimum based on policy settings for balance and consolidation, yet be almost out of resources. It is important to consider both measures when managing your environment.

### How Does VMware Aria Operations Calculate and Forecast Capacity

Calculating and Forecasting Capacity
Capacity analytics helps you assess the utilization and capacity remaining in objects across your environment. An evaluation of the historical utilization of resources generates a projection of the future workload. You can plan for infrastructure procurement or migrations based on the projection and avoid the risk of capacity shortage and high infrastructure costs.

Capacity analytics uses the capacity engine to assess historical trends, which include utilization peaks. The engine chooses an appropriate projection model to predict the future workload. The amount of historical data that is considered depends on the amount of historical utilization data.

#### Capacity Engine and Calculations

The capacity engine analyzes historical utilization and projects future workload by using real-time predictive capacity analytics, which is based on an industry-standard statistical analysis model of demand behavior. The engine takes the Demand and Usable Capacity metrics as input and generates the output metrics, which are Time Remaining, Capacity Remaining, Recommended Size, and Recommended Total Capacity, as shown in the following figure.

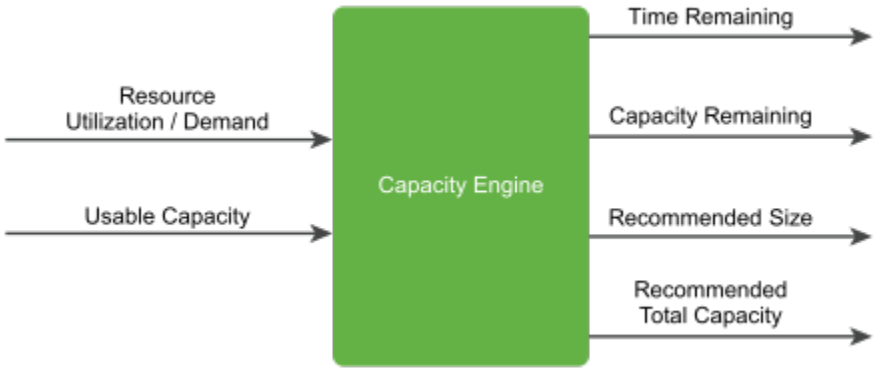


The projection window for the capacity engine is 1 year into the future. The engine consumes data points every 5 minutes to ensure real-time calculation of output metrics.

The capacity engine projects the future workload in a projected utilization range. The range includes an upper bound projection and a lower bound projection. Capacity calculations are based on the time remaining and risk level. The engine considers the upper bound projection for a conservative risk level and the mean of the upper bound projection and lower bound projection for an aggressive risk level. For more information about setting risk levels, see, *`Capacity Details`* in the Configuring Policies chapter of the VMware Aria Operations Configuration Guide.

The capacity engine calculates the time remaining, capacity remaining, recommended size, and recommended total capacity.

You can define custom metrics to compute capacity. The custom metric configuration is available in the policy settings. The calculation algorithm itself does not change. The custom metrics that you configure in the policy is used by VMware Aria Operations in the capacity calculations. You can select the default metrics shipped with VMware Aria Operations, or

create super metrics and select them for custom capacity calculations. Changes that you make take effect after the next collection cycle.

**NOTE**
Enabling custom metrics in the capacity calculations is an advanced configuration. Custom metrics alter the way VMware Aria Operations calculates capacity across your enivronment. Use this setting only when needed.

For more details on how to enable custom metrics, see the topic, Capacity Details.

**Time Remaining**

The number of days remaining till the projected utilization crosses the threshold for the usable capacity. The usable capacity is the total capacity excluding the HA settings.

**Capacity Remaining**

The largest difference between the usable capacity and the projected utilization between now and 3 days into the future. If the projected utilization is above 100% of the usable capacity, the capacity remaining is 0.

**Recommended Size**

The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The warning threshold is the period during which the time remaining is green. The recommended size excludes HA settings.

If the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization 150 days into the future.

VMware Aria OperationsVMware Cloud Foundation Operations caps the recommended size that is generated by the capacity engine to keep the recommendations conservative.

- VMware Aria OperationsVMware Cloud Foundation Operations caps an oversized recommended size at 50% of the currently allocated resources.
  For example, a virtual machine that is configured with 8 vCPUs has never used more than 10% CPU historically. Instead of recommending a reclaim of 7 vCPUs, the recommendation is capped to reclaiming 4 vCPUs.
- VMware Aria OperationsVMware Cloud Foundation Operations caps an undersized recommended size at 100% of the currently allocated resources.
  For example, a virtual machine that is configured with 4 vCPUs has been constantly running very hot historically. Instead of recommending the addition of 8 vCPUs, the recommendation is capped at adding 4 vCPUs.

**Recommended Total Capacity**

The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The recommended total capacity includes HA settings.

For example, if the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization including HA values, 150 days into the future.

**NOTE**
Recommended total capacity is not available for objects.

The following figure shows the capacity calculations for a conservative risk level.

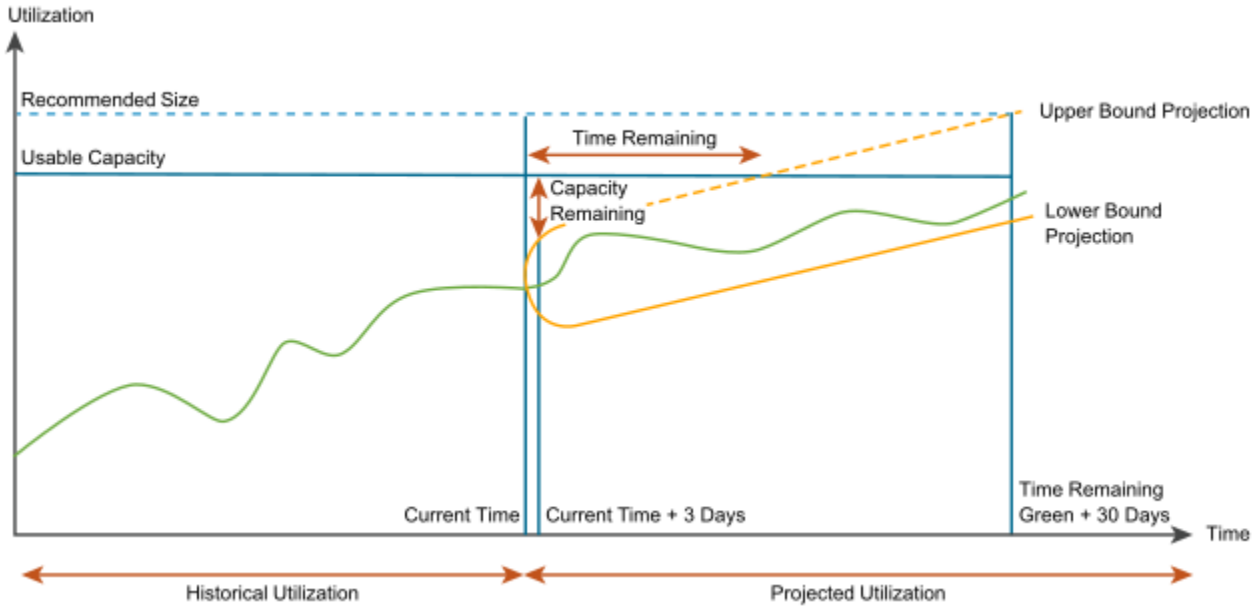


The following figure shows the capacity calculations for an aggressive risk level.

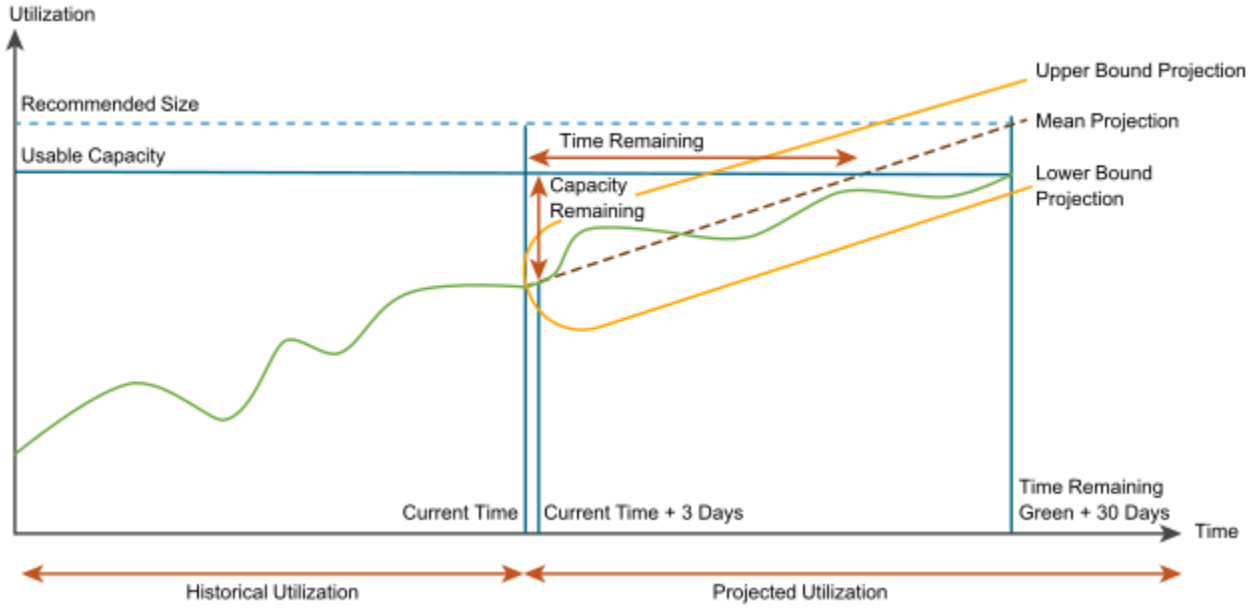

**NOTE**

- If HA is not enabled in VC then Usable Capacity = Total Capacity. In this case, the Usable Capacity value can be 0 only if there are no hosts in the Cluster.
- If HA is enabled, then Usable Capacity can be 0 in the following cases:
  - There are no hosts in the cluster.
  - HA is configured incorrectly. For example: it can be configured to 100% percent. Please check the HA configuration in vCenter.
  - HA Active host count is less than 2.
  - The host is not HA Active if:
    - Host is in Maintenance Mode.
    - Host is Powered Off.
    - Value of “runtime.dasHostState” property is not equal to “connectedToMaster” or “master”. This can be because of some network issues between the hosts.

## Utilization Peaks

The historical utilization of resources can have peaks, which are periods of maximum utilization. The projection of future workload depends on the types of peaks. According to the frequency of peaks, they can be momentary, sustained, or periodic.

**Momentary Peaks**

Short-lived peaks that are a one-time occurrence. The peaks are not significant enough to require additional capacity, so they do not impact capacity planning and projection.

**Sustained Peaks**

Peaks that last for a longer time and impact projections. If a sustained peak is not periodic, the impact on the projection lessens over time because of exponential decay.

**Periodic Peaks**

Peaks that exhibit cyclical patterns or waves. The peaks can be hourly, daily, weekly, monthly, during the last day of the month, and so on. The capacity engine also detects multiple overlapping cyclical patterns.

## Projection Models

The capacity engine uses projection models to generate projections. The engine constantly modifies projections and chooses the model that best fits the pattern of historical data. The projection range predicts the general usage pattern that covers 90% of the future data points. Projection models can be linear or periodic.

**Linear Models**

Models that have a steadily increasing or decreasing trend. Multiple linear models run in parallel and the capacity engine chooses the best model.

Examples of linear models are linear regression and autoregressive moving average (ARMA).

**Periodic Models**

Models that discover periodicity of various lengths, such as hours, days, weeks, months, or the last day of the week or month. Periodic models detect square waves that represent batch jobs and handle data streams that contain multiple overlapping periodic patterns. These models ignore random noise.

Examples of periodic models are fast Fourier transforms (FFTs), pulses (edge detection), and wavelets.

## Forecast In Trend Views

Forecasts are generated based on the time range specified in the view settings and are forecasted for the number of days specified in the forecast setting. The forecast is generated based on 3 main algorithms. Change-point detection to find

sections of the history with significant changes, linear regression to find linear trends, and cyclical analysis to identify periodic patterns.

### Historical Data Window

The capacity engine captures historical data over a period of time depending on the historical data window. The historical data window that the engine uses is an exponential decay window.

The exponential decay window is a window of unlimited size in which the capacity engine gives more importance to the most recent data points. Beginning from the projection calculation start point, the engine consumes all the historical data points and weighs them exponentially, based on how far back in time they are.

## Allocation and Demand Model in Workload Optimization

VMware Aria Operations uses the demand model and allocation model for Workload Optimization. By default, only the demand model is used. You can turn the allocation model on in the active policy.

### What is Allocation Model?

The allocation model determines how much compute, memory, and storage resources are allocated to object types. You define the allocation values by modifying the policy which is applied to the objects. The allocation values, also known as overcommit ratios, affect performance and cost.

The allocation model works alongside the demand model. Unlike the demand model which always affects the capacity calculations, the allocation model can be turned on or off in the policy setting. You can control the ratio by which VMware Aria OperationsVMware Cloud Foundation Operations overcommits either the CPU, memory, or disk space. By specifying the allocation values in the policy, you can choose whether you want to overcommit your resources or not. Overcommitting helps you measure utilization of resources in a pay-as-you-go model. When you do not overcommit, the utilization of your cluster will never exceed 100%. If your resource utilization is over the allocation ratio that you set, Capacity Remaining becomes zero.

To modify a policy and configure overcommit ratios, see Policy Allocation Model Element in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.*

### What is Demand Model?

The demand model is a concept used to predict future resource requirements based on historical data and usage patterns. The demand model is a key component of capacity planning and optimization, allowing administrators to anticipate the growth of workloads and infrastructure resources needed to support those workloads over time.

The demand model takes into account various factors, such as historical resource consumption, trends, seasonality, and potential changes in workloads. By analyzing this information, VMware Aria Operations can provide insights into when additional resources (such as CPU, memory, storage, and network capacity) will be needed to meet the expected growth in demand.

### Difference between Allocation Model and Demand Model in Workload Optimization

When clusters in a data center are not optimized, VMware Aria Operations uses the demand model by default to optimize the VMs placed in the cluster. While VMware Aria Operations moves VMs to optimize the cluster, it uses the capacity demand to calculate the best way to optimize the cluster.

When you turn on the Allocation Model in the active policy, VMware Aria Operations optimizes based on the Demand Model, and takes into consideration the overcommit ratios defined by the Allocation Model. As the VMware Aria Operations WLP engine takes into the consideration the overcommit ratios based on the allocation model for CPU, memory and disk space, these resources will not be under stress after optimization due to the over allocation. In any optimization model which is based on allocation the target is always to reduce the stress on all clusters and ensure that after the VM migration the target clusters are not over-allocated.

**How to Enable Allocation Model**

From the left menu, click the **Policy Definition** tile under **Operations › Configurations**. Select the active policy and click the **Capacity** card. In the Capacity Setting section, click the lock icon next to the Allocation Model for your vCenter object type. The available object types are:

- Cluster compute resource. Set the overcommit ratios for CPU, memory and disk space. You can also consider powered off VMs by selecting the **Activate** checkbox.
- Datastore. Set the overcommit ratio for the vSAN datastore disk space.
- Datastore cluster. Set the overcommit ratio for cluster disk space.

After you have set the overcommit ratios, click **Save**. After setting the overcommit ratios for the Allocation Model in the active policy, go to the Workload Placement page to see the Allocation Model displayed under the **Are your clusters meeting your utilization objective?** section.

### Policy Settings for Capacity

VMware Cloud Foundation Operations lets you define capacity and workload automation policies on objects and object groups. These policies help you define custom metrics, and set symptom thresholds, alerts, and recommendations for capacity planning.

To set the capacity settings for your policy, from the left menu, **Operations › Configurations** and then click **Add** in the **Policy Definition** tile. Or select the required policy and edit it.

For more information on policy settings for capacity, see the topics, Capacity Details and Workload Automation Details.

## How to View and Assess Capacity

VMware Cloud Foundation Operations helps you asses capacity by predicting future resource needs based on historical data and usage patterns. VMware Cloud Foundation Operations helps you analyze trends, forecast demand, and allocate resources accordingly to prevent capacity shortages and overprovisioning. To view capacity, you can open the Capacity page or view the Capacity Tab in the Inventpry page.

### Assessing Capacity in the Capacity Page

The Capacity page in VMware Cloud Foundation Operations helps you assess workload status and how much capacity is remaining in data centers across your environment.

**Where You Find the Capacity Page**

Click **Assess** under **Capacity** in the left menu.

> **NOTE**
> Click on a data center graphic to display the object details screen for the data center.

For more details on how to use the capacity page to optimize capacity, see the topic, Using the Capacity Page to Asses and Optimize Capacity.

For more details on viewing objects on the capacity page, see the topic, Viewing Object Capacity in the Capacity Tab.

### Viewing Object Capacity in the Capacity Tab

The **Capacity** tab provides Time Remaining and Capacity Remaining data for the selected object. Virtual Machine Remaining data is the sum of Virtual Machine remaining from each Cluster Compute Resource and is available for

datastores, datastore clusters, clusters, data centers, CDC, and VC based on the average profile, or when you activate one or more custom profiles in the policy.

### Where You Find the Capacity Tab at an Object Level

From the left menu, click **Inventory**, then select a group, custom data center, application, or inventory object. Click the **Capacity** tab under **Details**.

### Understanding the Capacity Tab

For the selected object, the **Capacity** tab lists two panes with the Time Remaining and Capacity information. These panes display the value of the resources remaining till they run out.

Below the **Time Remaining**, **Capacity** and **Virtual Machine Remaining** panes, the time and capacity utilization metric for CPU, memory, and disk space are displayed in three panes. By default, the most constrained resource is selected. Click **CPU**, **Memory**, or **Disk Space** to change the views to these resources. These panes display the resource information based on the Demand model (default) or Allocation model (if configured).

For datastores, if you have activated Allocation Model and Capacity buffer in the assigned profile, you see the **Disk Space** information based on Allocation and Usage.

#### Time Remaining Pane

When you select the **Time Remaining** pane and click one of the resource types, the utilization graph displays the historical value of the utilization metric and its forecast plotted against time, projecting how swiftly resource utilization is approaching the usable capacity.

Click **RESET** in the **Time Remaining** pane if you want to change the date from when the historical utilization is calculated. By default, it is calculated from the object creation or VMware Aria OperationsVMware Cloud Foundation Operations installation date. Click **RESET** if you want to change projected utilization, for example drop the irrelevant historical data from the calculation.

This affects the capacity calculation for the future trend, which will impact Capacity Remaining, Time Remaining, Recommended Size, and VM Remaining (if it is available for the particular resource, namely, Datastore, Cluster, Datastore Cluster). It will also impact all the resource containers (CPU, Memory, Disk Space, etc.)

For example, if the capacity calculation is using historical data based on the duration when you were provisioning a large number of VMs, but in the recent past, there was no provisioning done. Then, VMware Aria OperationsVMware Cloud Foundation Operations may not project a trend that is based on the recent data when there was no provisioning, but still will consider historical data when there were a large number of VMs provisioned. This historical data may show an increasing trend. In this case, you can change the capacity computation to start from a date after the VM provisioning is complete.

> **NOTE**
> If you are unable to see the **RESET** button in the **Time Remaining** pane, make sure you have the **Manage Capacity Calculation** permission under **Administration › Control Panel › Access Control › Roles › Capacity › Assess**.

#### Capacity Pane

The **Capacity Remaining** pane indicates the unused capacity of your virtual environment to accommodate new virtual machines. VMware Aria OperationsVMware Cloud Foundation Operations calculates the Capacity Remaining as a percentage of the remaining capacity, compared to the total. Capacity Remaining is calculated as the utilization metric forecast 3 days from now subtracted from the Usable Capacity. VMware Aria OperationsVMware Cloud Foundation Operations calculates the average profile and always computes the virtual machine remaining number based on the average profile. You can change the profile by clicking the + icon above the bar chart. VMware Aria OperationsVMware Cloud Foundation Operations calculates virtual machine remaining numbers when you activate one or more custom profiles from the policy. The overall virtual machine remaining is based on the most constrained profile.

When you select Capacity and click one of the resource types, a bar chart and a table of values based on the Demand and Allocation model (if configured) appears. The bar chart displays total usable resource, the percentage used, the percentage allocated for high availability and buffer, and the percentage remaining based on the Demand and Allocation models (if configured).

The table displays the following information for each resource type:

- **Total**: The total usable capacity for each resource type based on the Demand model or Allocation model (if configured). The difference in Total capacity and Usable capacity is set in the HA (admission control) that is set in the clusters in vSphere.
- **Usable**: The total usable capacity for each resource type based on the Demand model or Allocation model (if configured).
- **Used**: Approximate value how much utilization do you have now. Shows the forecast value of utilization metric in 3 days from now.
- **Recommended Size**: The Total Capacity that must be available for a green level of Time Remaining. The slider in the policy controls the Time Remaining green zone, and the default value is 150 days.
- **Remaining**: The Capacity Remaining metric value and also the percentage. The value of Capacity Remaining metric is calculated by forecasting the utilization metric 3 days from now and subtracting it from Usable capacity.
- **Buffer**: The percentage of the capacity buffer based on the buffer value that you set in the policy. The Capacity Buffer element determines how much extra headroom you have and ensures that you have extra space for growth inside the cluster when required.
- **High Availability**: The percentage of the high availability based on the high availability buffer.

**Virtual Machine Remaining**

The virtual machine remaining number is based on the average profile. The virtual machine remaining numbers are calculated when you activate one or more custom profiles from the policy. The overall virtual machine remaining is based on the most constrained profile.

When you click on Virtual Machine Remaining, you see the number of virtual machines based on the Average Profile. Select a different policy by clicking **Assign Policy** from the drop-down list on the top right hand side of the page. To activate custom profiles in that policy, or other policies, click the **GO TO CUSTOM PROFILES** link. For more information, see the topic, *Custom Profiles in VMware Aria OperationsVMware Cloud Foundation Operations*.

Only vSAN datastore objects, and not other datastore objects, display disk space based on Buffer and HA.

The **Capacity** tab is a subset of the Capacity optimization capability. For additional details, refer to Assessing Capacity in the Capacity Page.

## Capacity VM Shortfall

VMware Aria OperationsVMware Cloud Foundation Operations has a new metric known as VM shortfall, the VM shortfall value is always positive. The metric counts all the negative VMs Remaining and then turns them into positive. The VM shortfall metric will be available for Clusters, Datastores, and Datastore Cluster objects only. The VM Shortfall metric details is displayed in the Capacity Summary pages next to VM Remaining, in case if VM Remaining is equal to zero.

For all these objects you can create custom profiles directly from the capacity page. In the **Capacity** tab, click the + sign under Virtual Machine remaining. From the **Applicable Profiles** page, click **Add** to add a new profile. You can either add new profile or import from existing object. You can modify the required metric value and click **OK**.

## How Does VM Shortfall Work

When the cluster has more VMs than its usable capacity can handle, the shortfall of capacity is shown by the VM counter. VM remaining will show zero and not negative. The shortfall appears only when the usable capacity is zero.

Whenever there is no capacity in that case, VM shortfall detects the VM profiles in the setup and finds the profile for which we have the maximum number of VMs that are overcommitted.

VM Shortfall Formula

It is possible that when one datastore is connected with multiple clusters, each cluster thinks that the whole diskspace available in the datastore is for itself. We can calculate the VM shortfall as:

```
VM Shortfall = (Used - Usable Capacity) / Profile Sizes
```

> **NOTE**
> In the above formula the Used Capacity represents the specific value that is displayed in the **Capacity › Capacity Remaining** tab.

If you have two clusters, one cluster has 20 VMs remaining and the other cluster has shortfall of 20 VMs, then the number of VMs remaining for the datacenter will be 20 VMs, however, if you have a shortfall of VMs in the datastore connected to that cluster, then we will see on the datacenter the remaining VMs as zero.

Metrics

The VM shortfall displays metrics under Capacity Analytics Generated and under Profiles. The metrics are:
- Capacity Remaining Profile (Average)
- Capacity Shortfall Profile (Average)

## vSAN HCI Mesh

The vSAN HCI Mesh is applicable if you have a vSAN OSA HCI cluster. The vSAN HCI Mesh allows vSAN clusters to remotely mount the datastore of another (remote) vSAN cluster, hence sharing the storage capacity and span its usage to a wider pool of compute resources. HCI Mesh allows multiple vSAN clusters to share their datastores remotely. If the vCenter version is 7.0U1 and greater you can share other cluster datastores and provision VMs to the remotely shared data stores.

A new card is created under **Capacity › Asses › vSAN HCI Mesh**. The vSAN HCI Mesh card is displayed only when there are eligible clusters which can be part of the vSAN HCI mesh. If the version of the vCenter is less than 7.0U1 for the selected datacenter, this card is not displayed. If the mesh is not configured, click the link, which lets you configure the clusters from vCenter.

The vSAN HCI Mesh displays the following information:
- Total Capacity - Displays the total capacity of all the vSAN datastores connected to the mesh.
- vSAN Mesh - Displays how the datastores are interconnected in the HCI mesh.
- Inner Circle - Represents the local datastore and displays capacity details like, free capacity, used capacity, and total capacity.
- Outer Circle - Represents the remote datastore and displays capacity details like, total mounted capacity, free capacity, and total capacity usage percentage.
- Arrow Heads - The arrow heads points to the clients connected to the server vSAN cluster.

**Datastore View**

Displays the capacity consumption details for the selected datastore. You can also identify which are the client and server clusters.

**Cluster View**

Displays the capacity details of the remote datastores. You can also view the capacity details for the mounted datastores. The Total Usage % displays the total capacity consumed by all the client clusters connected to these mounted datastores including the selected cluster.

**Health Alerts Related to HCI Mesh**

Displays the health alerts related to the HCI Mesh. You can also view alerts for individual datastores, which is part of the HCI mesh.

## vSAN Remote Datastores

The vSAN Remote Datastore card is displayed if you have a vSAN Max ESA cluster and a vSAN Compute cluster mounting the vSAN Max cluster. The vSAN Compute cluster does not have any storage of its own and uses the storage provided by the vSAN Max cluster. vSAN Max cluster allows vSAN compute to remotely mount it's datastore, hence sharing the storage capacity and span its usage to a wider pool of compute resources. vSAN Max cluster allows multiple vSAN Compute clusters to share their datastores remotely.

The vSAN Remote Datastores card displays the following information:

- Total number of clusters mounted: Displays the total number of the vSAN Compute clusters mounted on the vSAN Max cluster.
- Total Capacity: Displays the total capacity of the vSAN datastores connected to the vSAN Max cluster.
- Free Capacity: Displays the remaining capacity of the vSAN datastores connected to the vSAN Max cluster.

Click **View Remote Datastores** for a more detailed view of the information.

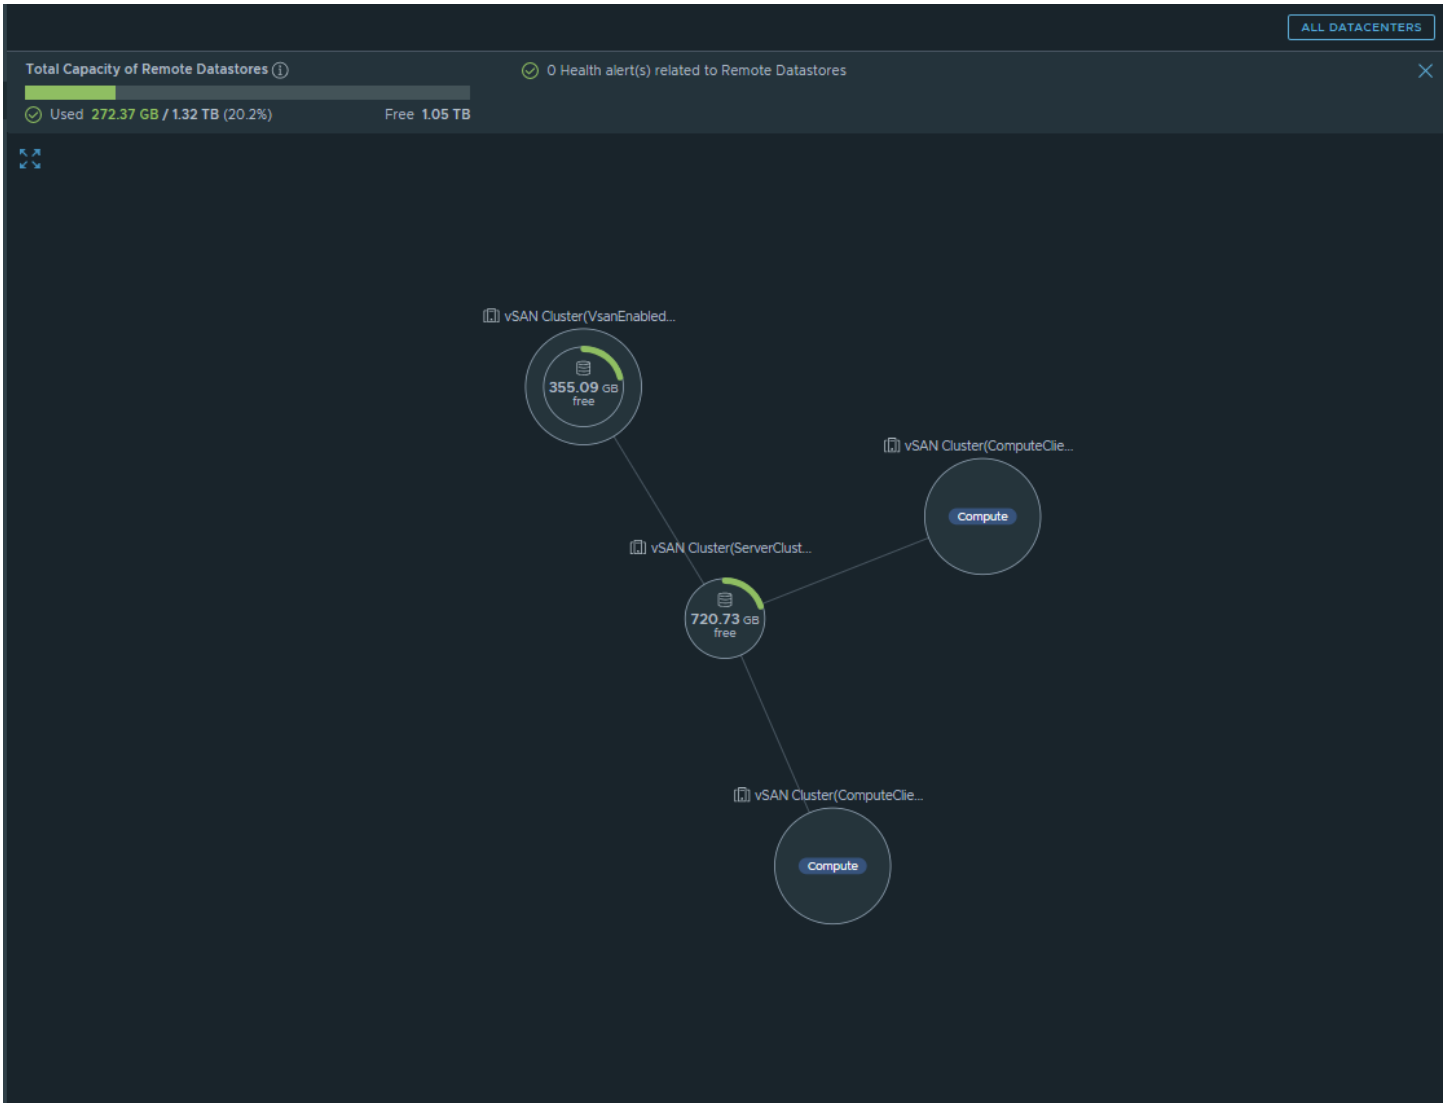


The total capacity details of the remote datastore and the health alerts (if any) of the remote datastore are displayed along with the relationship between the vSAN Compute clusters and the remote datastore they use.

You can view which vSAN Compute cluster is mounting the vSAN Max cluster and their usage trends. Hover over each cluster type to view the details related to the it. For example, if you hover over a vSAN Compute cluster, you can see the

number of datastore it is using along with the total amount of mounted and free capacity in gigabytes and the total usage percentage of the cluster. If you hover over the vSAN Max cluster, you can view the total, free, and used capacity in gigabytes, and the usage percentage of the datastore.

**Custom Profiles in VMware Aria OperationsVMware Cloud Foundation Operations**

Custom Profiles
Custom Profiles in VMware Aria OperationsVMware Cloud Foundation Operations
A custom profile defines a hypothetical configuration of a virtual machine. Custom profiles help you determine how many instances of that virtual machine can fit in your environment, depending on the capacity remaining and the configuration of the parent object. You activate the custom profile in policies for specific object types. The custom profiles are used to calculate VM remaining when the policy is selected for those object types in the **Capacity** tab of the VMware Aria OperationsVMware Cloud Foundation Operations object browser.

## How Custom Profiles Work

As with default profiles, custom profiles is based on metric values for a virtual machine. You can add as many custom profiles as you need for an object. For example, you might create one custom profile for a virtual machine that has a memory demand model of 2 GB. You create another custom profile that has a memory demand model of 4 GB.

VMware Aria OperationsVMware Cloud Foundation Operations uses custom profiles to calculate the number of instances of those virtual machines that can fit in your environment for specific objects. The number of virtual machines is based on the capacity allocation and demand defined for the object.

After you define the metric values for a custom profile, you configure for which policies and object types the custom profiles are activated.

## How Custom Profiles are Related to Policies

After you create a custom profile, you activate it in a policy for one or more of the following object types:
- Datastores
- Datastore clusters
- Cluster compute resource

The custom profile is then available for all objects of that object type in the **Capacity** tab of the VMware Aria OperationsVMware Cloud Foundation Operations object browser.

To determine how many instances of the virtual machine with a specific configuration as defined by the custom profile you can include in the parent object, go to the **Capacity** tab of the parent object. Select the **Capacity Remaining | Virtual Machine Remaining** pane. Select the policy in which the custom profile is activated, from the top right hand side drop-down list. The custom profiles appear in the VM remaining section and indicate how many instances of the virtual machine can fit in your environment. The overall virtual machine remaining is based on the most constrained profile.

## Where You Find Custom Profiles

Click the **Custom Profiles** tile under **Operations › Configurations**. A list of all the custom profiles available in VMware Aria OperationsVMware Cloud Foundation Operations is displayed in a table. Click the name of a custom profile from the **Name** column to view the metric values and policies of that custom proile on the right pane.

**Table 113: Options in the Custom Profiles Page**

| Option | Description |
|---|---|
| Toolbar options | In the toolbar click **Add** to add a custom profile for a specific object type. Click the **Vertical Ellipses** against a profile to perform the following actions:<br>• **Delete**. Delete the selected profile.<br>• **Export**. Export the selected profile.<br>• **Import**. Import selected profiles. |
| Sorting options | Filter the list to display profiles that match the filter you create. You can sort by name, description, object type, or adapter type. |
| Filter and Search | Use the advanced filter and search box to look for custom profiles. Click the drop-down icon to search by name, description, object type, adapter type. |

### Custom Profiles Add and Edit Workspace

You can add a custom profile for a cluster compute resource, datastore cluster, or datastore, to determine how many VMs of the specific object type can fit in your environment. In the Custom Profiles workspace, you create a custom profile for the object and define the metrics for the virtual machines.

#### Where You Create or Edit a Custom Profile

To create a custom profile, click the **Custom Profiles** tile under **Operations › Configurations**. The Custom Profiles page displays all the available custom profiles in a table. You have the option to add custom profiles via the **ADD** button, or filter and search for profiles using the search box. To see the properties of a custom profile, click the custom profile name in the table. The custom profile details are displayed on the right hand side pane.

Click the **Vertical Ellipses** next to a profile name to edit, delete, clone, or export the profile. Alternatively, to delete or export the more than one profile, select the checkbox next to the profile name and click the **Vertical Ellipses** next to the **ADD** button, and perform an action. You can also import custom profiles after clicking the **Vertical Ellipses** next to the **ADD** button.

#### Add or Edit Custom Profiles Wizard

When you click the **ADD** button, or **Edit** next to the **Vertical Ellipses**, you get a three step wizard. Click the **Next** button after every step to proceed with the wizard.

**Table 114: Profile Step**

| Option | Description |
|---|---|
| Name | Descriptive name of the custom profile. This field is mandatory. |
| Description | Meaningful description for the custom profile. Provide specific information that other users must know about this profile. |
| Profile Type | The profile type for which the custom profile is created. Currently, you can create a custom profile only for a virtual machine. The Capacity Remaining is calculated for this object type. |

**Table 115: Metrics Step**

| Option | Description |
|---|---|
| Define Values for Profile | Populate the value and unit for the capacity metrics. You can also copy the metric values for the custom profile from an existing object. When you click the **Copy from Existing Object** button, the **Select an Object** pane opens on the right hand side. Browse to the object whose metric values you want to copy, or use the advanced search and filter to search by VM Name, vCenter, VM Tag, or Custom Group. Click **Copy** when done. |

**Table 116: Policies Step**

| Option | Description |
|---|---|
| Checkbox | Select the checkbox against the name of the policy which is displayed in rows for one or more of the following types of objects:<br>• Cluster Compute Resource<br>• Datastore<br>• Datastore Cluster |

Click **CREATE** to create the custom profile or **UPDATE** if you are updating an existing custom profile.

### Clone a Custom Profile

Click the **Vertical Ellipses** next to a profile name to clone an existing profile profile. When you click **Clone**, you get a three step wizard. In the **Profile** step of the wizard, you can provide the name, description, and select the profile type. In the **Metrics** and **Policies** steps of the wizard, the details from the custom profile that you are cloning from, are already populated. You can skip these steps and click the **CREATE** button to create a cloned custom profile with the same values and policies as the original custom profile.

### Custom Profiles Import and Export

Custom Profiles Import and Export
Import or Export Custom Profiles in VMware Aria OperationsVMware Cloud Foundation Operations
Export a custom profile as an XML file and import it again in a different VMware Aria OperationsVMware Cloud Foundation Operations instance to easily transfer custom profile configurations.

### Where you Import or Export Custom Profile

Navigate to the **Custom Profiles** tile under **Operations › Configurations**.. Do one of the following:

- To import a custom profile, click the Vertical Ellipses next to the **ADD** button and select Import. The **Import Custom Profile** dialog box opens. Browse to the XML configuration file and select what VMware Aria OperationsVMware Cloud Foundation Operations has to do in case of a conflict. Click **IMPORT** to complete the process.
- To export an existing custom profile, select the box next to the custom profile name. Click the Vertical Ellipses next to the name and select **EXPORT**. The XML configuration file is downloaded to you local system via the web browser.

### Predefined Dashboards for Capacity

VMware Cloud Foundation Operations provides predefined dashboards to help you monitor and manage your virtualized environments. Capacity dashboards help you analyze the capacity of your environment by providing insights into resource consumption trends, utilization patterns, and potential bottlenecks. They can help you plan for future growth and optimize resource allocation.

For more information on the predefined dashboards available for capacity optimization in VMware Cloud Foundation Operations, see the topic, Capacity Dashboards.

## How to Optimize Capacity and Improve Performance in VMware Cloud Foundation Operations

How to Optimize Capacity and Improve Performance
VMware Cloud Foundation Operations helps you optimize capacity by identifying and reclaiming unused or underutilized resources, such as reclaiming idle CPU or memory from VMs. This practice can help free up resources for other workloads. Rightsizing helps you adjust the resources allocated to virtual machines based on their actual usage. v
VMware Cloud Foundation Operations identifies overprovisioned or underutilized VMs and provides recommendations to optimize their resource allocations.

### Using the Capacity Page to Asses and Optimize Capacity

The capacity page in VMware Aria Operations offers a comprehensive view of resource utilization, allocation, trends, and forecasts, allowing administrators to make informed decisions about resource provisioning, performance optimization, and future growth.

To open the Capacity page, click **Assess** under **Capacity** in the left menu.

#### How Optimize Capacity Works

The Capacity Optimization and Reclaim features are tightly integrated functions that allows you to assess workload status in data centers across your environment. You can determine time remaining until CPU, memory, or disk space resources run out and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

When you open the Capacity page, graphical representations of all the data centers and custom data centers in your environment appear. VMware Cloud on AWS data centers has a unique icon to differentiate it from the other data centers.

By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To review the status of a data center, click the graphic. The page refreshes to display the following data:

**Time Remaining**

Time Remaining specifies which clusters are most constrained and displays the criticality of the cluster.

**Optimization Recommendations**

VMware Aria OperationsVMware Cloud Foundation Operations shows you the number of reclaimable VMs and the associated cost savings. Click **View Reclaimable VMs** to navigate to the **Reclaim** page.

**Cluster Utilization**

Cluster Utilization displays an interactive graph that shows time remaining by component. You can explore the demand percentage over time by CPU, memory, and disk space or by the most constrained component. By default, the data displayed is for the Demand model. If you have configured the Allocation model, then you can also see the CPU, memory, and disk space time remaining model based on the overcommit ratios that you have set in the policy.

Click the **Edit** icon to modify the criticality threshold, risk level, and allocation model. These changes affect the selected cluster's policy. Hence, any change that you make here, affects all the clusters under the same policy.

Set the **Show History** and **Show Forecast** variables to create the slice of time in which you want to see time-remaining data. The vertical axis of the graph shows the total capacity being used by the current amount of CPU, memory, or disk space respectively. The bold, black line across the top of the graph depicts the historical value of usable capacity. The horizontal axis is the timeline. Vertical lines in the graph are labeled at the bottom of each line. The first vertical dotted line on the left marks the projection calculation start point. The next line is the current date - now. The third vertical marks the date the resource runs out. If a resource has little time remaining, the current date and the date that time runs out may be the same.

VMware Aria OperationsVMware Cloud Foundation Operations can make recommendations for increasing time remaining based on the data it receives and these recommendations appear at the bottom of the screen. You might see two options: Option 1 shows what you can achieve by reclaiming resources. Option 2 shows the results of adding capacity.

If you choose to reclaim resources, you can run that process immediately by clicking **RECLAIM RESOURCES**. To see the details or choose additional options before running a reclaim action, review the information provided in the **Optimization Recommendations** pane and then click **VIEW RECLAIMABLE VMS** to go to the **Reclaim** page.

**Table 117: Capacity Optimization Options**

| Option | Description |
|---|---|
| **Select a datacenter** | Select a data center from the carousel across the top of the page. Information about the datacenter is displayed below. |
| **ALL DATACENTERS \| X** | Toggle: click **ALL DATACENTERS** on the upper right when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers. |
| View: | Filter results to include data centers, custom data centers, or both. This option appears if you select **ALL DATACENTERS** on the upper right. |
| Group BY: | Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. This option appears if you select **ALL DATACENTERS** on the upper right. |
| Sort by: | Options (Options appear if you select **ALL DATACENTERS** on the upper right):<br>• Alarm clock graphic - lists data centers/custom data centers by time remaining.<br>• Dollar sign - lists data centers/custom data centers by potential cost savings.<br>• Scales graphic - lists data centers/custom data centers by level of optimization. |
| **Select datacenter or ADD NEW CUSTOM DATACENTER** | Options (options appear if you select **ALL DATACENTERS** on the upper right):<br>• Select a data center from the carousel across the top of the page. All data following refreshes with information for the selected object.<br>• Select **ADD NEW CUSTOM DATACENTER** to display a dialog box that allows you to define a custom data center. |
| Time Remaining | Appears when you select a data center or custom data center from the top of the screen.<br>Gives overview of cluster status, including how many are at:<br>• Critical |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Medium<br>• Normal<br>• Unknown<br><br>"Critical" can indicate a resource contention, imbalance, or other stress condition. Thresholds you set in the policies define what is critical. |
| Optimization Recommendations | Lists potential cost savings by reclaiming unused resources.<br><br>Indicates if workloads can be optimized across clusters.<br><br>**VIEW RECLAIMABLE VMS** - displays the **Reclaim** screen, where you can research and run potential VM reclamation actions.<br><br>**VIEW OPTIMIZATION** - displays the **Workload Optimization** screen, where you can optimize workloads based on your policy settings. |
| **Cluster Utilization and Time Remaining** | Overall view of cluster health in the selected data center. You can select a cluster from the list to display information about that cluster, or use the options to sort and filter results. The options you select dictate the data displayed in the graph.<br><br>Sort by:<br><br>• Most Constrained: most constrained element<br>• CPU (allocation or demand)<br>• Memory (allocation or demand)<br>• Disk Space (allocation or demand)<br><br>**NOTE**<br>Demand model is always on and is the default.<br><br>**Filter**: search field.<br><br>**Show History for**: The period before forecasting begins (does not impact the forecast calculation).<br><br>**Show Forecast For**: The forecast period.<br><br>**How is the criticality determined?** Displays the criticality threshold you set for this type of object in the Policies Library.<br><br>**Cluster Time Remaining Settings**: Click the **Edit** icon to edit the default policy for the selected cluster. Change the criticality threshold, risk level, allocation model and capacity buffer. Applying these changes affects all objects in the policy. For more information, see Configuring Policies in the *VMware Aria Operations Configuration Guide* *VMware Aria Operations SaaS Configuration Guide* |
| Time Remaining graph | Data shows current and trending resource usage and pinpoints when a given cluster is projected to run out of CPU, memory, or disk space based on the allocation or demand model (default). VMware Aria OperationsVMwa |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | re Cloud Foundation Operations takes into account the business hours which you set in the current policy. |
| Legend | Displays a legend of the colours in the time remaining graph. Click on any one of the legends to toggle its display in the chart. |
| Recommendations | Option 1: Reclaim Resources.<br>Shows resources that can be reclaimed to increase time remaining for the selected cluster.<br><br>**RECLAIM RESOURCES** - displays the **Reclaim** screen, where you can research and run potential VM reclamation actions.<br><br>Option 2: Add Capacity.<br>Shows resources that can be added to increase time remaining. |

**NOTE**
You might see that a data center or cluster is labeled optimized when it has few or no days remaining before CPU, memory, or disk space is predicted to run out. The seemingly odd assessment is due to optimization and time remaining being two different measures of data center and cluster health. A data center can be running at optimum based on policy settings for balance and consolidation, yet be almost out of resources. It is important to consider both measures when managing your environment.

### Using Reclaim to Free Up Resources

Use the **Reclaim** feature in VMware Cloud Foundation Operations to identify underutilized workloads and reclaim resources from across your environment.

#### Where You Find Reclaim

Select **Reclaim** under **Capacity** in the left pane.

**NOTE**
Double-click on a data center graphic to display the object details screen for the data center.

#### How Reclaim Works

The Capacity Optimization and Reclaim features are tightly integrated functions that allows you to assess workload status and resource contention in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out, and realize cost savings when underutilized VMs can be reclaimed and deployed where needed.

When you open the **Reclaim** page, graphical representations of all the data centers and custom data centers in your environment appear. By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To review the status of a data center, click the graphic. The area following refreshes to display details about the selected data center. The **How much you can potentially save** pane reflects potential capacity savings and indicates a possible cost savings once you have reclaimed underused or powered off VMs. The **Total Reclaimable Capacity** pane gives details of the reclaimable percentages for CPU, memory, and disk space.

The table at the bottom of the page provides important information about the VMs that offer the most cost savings. The VMs are listed by **Powered Off VMs**, **Idle VMs**, **Snapshots**, and **Orphaned Disks**. The highest priority heading is at the far left. You can specify what information is included in your reclaim action. For example, when you click a column heading, the table lists, by data center and then by VM, the allocated and reclaimable CPUs and memory, respectively.

Then, for example, you can select the box next to one or more VM names and click the **EXCLUDE VM(S)** button to keep those VMs from being included in any reclaim action. You can also select VMs to resize.

### Reclamation Settings

Select the gear icon next to the page heading to customize Reclamation Settings. This affects all data centers. Using the Reclamation Settings, you can exclude, for example all snapshots from being included in the reclaim action - by deselecting the Snapshots check box. Similarly, you can include or exclude powered-off VMs, idle VMs, and orphaned disks. For more information, see Reclamation Settings.

> **NOTE**
> To provide read-only access to the Reclamation Settings page for a user, configure the user role in the Access Control page (Roles tab) under **Administration › Control Panel › Access Control**. After you click **Add** or after you edit an existing role, select the **Manage Global Settings** permissions under **Administration › Global Settings** in the **Assign Permissions** section to grant access to modify the Reclamation Settings page. Unselect the **Manage Global Settings** permissions to grant read-only access.

### Run a Reclaim Action

Run a reclaim action as follows:

1. In the table headings, **Select** the types of VMs to reclaim.
2. **Click** the name of a listed cluster to show its VM list.
3. **Select** each VM or snapshot you want to reclaim.
4. Click **Delete VM(s)** to reclaim their resources.

**Table 118: Reclaim Options**

| Option | Description |
|---|---|
| Select a data center. | Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object. |
| **ALL DATACENTERS \| X** | Toggle: click **ALL DATACENTERS** on the upper right when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers. |
| View: | Filter results to include data centers, custom data centers, or both. Option appears when you select **ALL DATACENTERS** on the upper right. |
| Group BY: | Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. Option appears when you select **ALL DATACENTERS** on the upper right. |
| Sort by: | Options (Options appear when you select **ALL DATACENTERS** on the upper right):<br>• Alarm clock graphic - list data centers/custom data centers by time remaining.<br>• Dollar sign - list data centers/custom data centers by potential cost savings.<br>• Scales graphic - list data centers/custom data centers by level of optimization. |
| **Select data center or ADD NEW CUSTOM DATACENTER.** | Options (Options appear when you select **ALL DATACENTERS** on the upper right): |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object.<br>• Select **ADD NEW CUSTOM DATACENTER** to display a dialog box that allows you to define a custom data center. |
| How much you can potentially save. | Appears when you select a data center or custom data center from the top of the screen. Shows the total calculated potential cost savings when you accept system reclamation recommendations. |
| Total Reclaimable Capacity | Lists potential cost savings for the selected data center when you reclaim unused resources.<br><br>Resource: CPU, memory, or disk space<br>Reclaimable Capacity: how much capacity is available to reclaim from idle resources<br><br>% Reclaimable: percentage of total CPU, memory, or storage you can reclaim. |
| Duration older than: | Shows idle or powered off VMs that have been idle or powered off for at least the selected time period: one week, two weeks, or a month. |
| Table of Potential Cost Savings | Tabular representation of the VMs, Idle VMs, Snapshots, and Orphaned disks in the selected data center from which resources can be reclaimed.<br><br>Click one of the elements - powered off VMs, idle VMs, and so on - to refresh the table with data for that element. The table lists the relevant clusters. To see the VMs hosted in a given cluster, click the chevron to the left of the cluster name.<br><br>Click the check box next to the VMs you want to act on, or click the check box next to the column heading VM Name to act on all the VMs.<br><br>Once you select a VM or VMs, the dimmed options above the table become visible, as follows.<br><br>**Exclude VM(s)**: The selected VMs are excluded from your subsequent action. Excluding VMs from a reclamation action can reduce the potential cost savings.<br><br>For powered Off VMs:<br>• **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more reclaim action for powered off VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>• **DELETE VM(s)**: Deletes the selected VMs.<br>• **EXCLUDE VM(s)**: Excludes the selected VMs.<br>• **EXPORT ALL**: Exports the list of powered off VMs into a CSV file.<br><br>For idle VMs: |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more reclaim action for idle VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>• **DELETE VM(s)**: Deletes the selected VMs.<br>• **POWER OFF**: Powers off the selected VMs.<br>• **EXCLUDE VM(s)**: Excludes the selected VMs.<br>• **EXPORT ALL**: Exports the list of idle VMs into a CSV file.<br><br>For Snapshots:<br>• **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more reclaim action for snapshots. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>• **DELETE SNAPSHOT(s)**: Deletes the selected snapshots.<br>• **EXCLUDE VM(s)**: Excludes the selected snapshot.<br>• **EXPORT ALL**: Exports the list of snapshots into a CSV file.<br><br>**SHOW\|HIDE EXCLUDED VMS**: Toggle displays or hides the list of VMs you previously excluded.<br><br>**NOTE**<br>By default, calculations for reclaimable resources are based on the demand model. But if you turn on the allocation model in the policy settings, the calculations are based on the allocation model.<br><br>For Orphaned Disks:<br>• **EXCLUDE DISK(S)**: Exclude the selected disks in the actionable list.<br>• **EXPORT ALL**: Exports the list of orphaned disks into a CSV file. You cannot reclaim orphaned disks from the UI. Instead, export the list into a CSV file and then reclaim the orphaned disks manually.<br><br>**NOTE**<br>VMware Aria OperationsVMware Cloud Foundation Operations reports orphaned VMDKs conservatively. There might be a false positive situation when the used VMDK is reported as orphaned, particularly if the VMDK is located on a datastore which is shared among multiple VCs, while not all the VCs are monitored by VMware Aria OperationsVMware Cloud Foundation Operations.<br><br>Check the accuracy of the VMDK reported as an orphaned disk, and then perform a reclamation. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **SHOW\|HIDE EXCLUDED DISKS**: Toggle displays or hides the list of disks you previously excluded. Excluded disks are not listed in the exported CSV file. |

**Reclamation Settings**

Displays information about powered off VMs, idle VMs, snapshots and orphaned disks. This information helps to identify the amount of resources that can be reclaimed and provisioned to other objects in your environment or amount of potential savings that can be done in each month.

The types of VMs are ranked in the order of their importance in a reclamation action. A VM whose attributes match more than one VM type is included with the higher-ranking VM type. Grouping the VMs this way eliminates duplicates during calculations. As an example, powered-off VMs are ranked higher than snapshots, so that a powered-off VM that also has a snapshot appears only in the powered-off VM group.

If you exclude a given type of VM, all VMs matching this type are included with the next lower-ranked group they match. For example, to list all snapshots regardless of whether their corresponding VMs are powered-off or idle, deselect the Powered-off VMs and Idle VMs check boxes.

Further, you can configure how long a given class of VMs must be in the designated state - powered-off, for example, or idle - to be included in the reclamation exercise. You also can choose to hide the cost savings calculation.

| Property | Description |
|---|---|
| Show Cost Savings | Controls whether to show Cost savings in the **Capacity** page. |
| VM Idleness Criteria | Defines the criteria to identify VMs that remain idle during the defined period of time.<br>You can set the following values:<br>• MHz of CPU consumption of VMs, and percentage of the time VMs have less than defined MHz of CPU within each day. |
| Show the following types of VMs for reclamation | Configure the types of VMs that are shown for reclamation. You can select the type of VMs which are shown for reclamation, and then configure paramters which those VMs have to meet to be eligible for reclamation. |
| Powerd Off VMs | VMs that have been continuously powered off during the defined period of time.<br>The total storage capacity used is reclaimable. Total storage reclaimable cost is computed by multiplying storage rate with storage utilization. The direct cost of VM is also attributed. |
| Idle VMs | You can configure the following parameters based on which VMware Aria OperationsVMware Cloud Foundation Operations calculates idle VMs:<br>• How many days the VM has been idle for.<br>• Number of days before which the VMs were provisioned based on which they are excluded. |

*Table continued on next page*

*Continued from previous page*

| Property | Description |
|---|---|
| Snapshots | VM snapshots that have existed for the entire defined period of time.<br>Snapshots of a VM use storage space and such storage is reclaimable. The reclaimable cost is computed by multiplying storage rate with reclaimable storage value. |
| Orphaned Disks | VMDKs on datastores that are not connected to any registered VMs and have not been modified during the defined period of time.<br><br>Orphaned disks are VMDKs which are associated with a VM which are not in inventory, but still available in a datastore. You can configure the minimum number of days for which VMDKs not related to any existing VM will be reported as orphaned and appear under Orphaned Disks in Reclaim page.<br><br>**NOTE**<br>You can navigate to**Cost/Price** section under **Administration › Global Settings** in the left menu, and change the value of the **Orphaned Disks Collection** time. At this time that you set, VMware Aria OperationsVMware Cloud Foundation Operations checks for orphaned VMDKs in vSphere Client instances. The settings for **Cost Calculation** and **Orphaned Disks Collection** are interrelated. The default value for **Cost Calculation** is 9:00 PM, and the default for **Orphaned Disks Collection** is 8:00 PM. It is recommended to schedule **Cost Calculation** after **Orphaned Disks Collection**. |

**NOTE**
If you are unable to make changes in the Reclamation Settings page, your user role in the Access Control page (Roles tab) under **Administration › Access Control** must be modified by an administrator. The **Manage Global Settings** permissions under **Administration › Global Settings** in the **Permissions** pane controls access to the Reclamation Settings page.

### Using Rightsize to Adjust Resource Allocation

Rightsizing in VMware Cloud Foundation Operations refers to the process of adjusting the resource allocations of virtual machines (VMs) to match their actual workload requirements. It involves optimizing the allocation of CPU, memory, storage, and other resources to ensure that VMs are neither overprovisioned (allocating more resources than necessary) nor underutilized (allocating fewer resources than needed). The goal of rightsizing is to achieve efficient resource utilization, improve performance, and potentially reduce costs Use the Rightsize page to alter the number of CPUs and amount of memory in oversized and undersized virtual machines.

#### Where You Find Rightsizing

Select **Rightsize** under **Capacity** in the left pane. Then, select a data center.

## How Rightsizing Works

The Capacity Optimization, Reclaim, and Rightsizing features are tightly integrated functions that allows you to assess workload status and resource usage in data centers across your environment. You can determine time remaining until CPU, memory, or storage resources run out, and realize cost savings when underutilized VMs can be reclaimed and deployed where needed. With this function, you can change CPU size and memory values for oversized and undersized virtual machines to achieve optimum system performance.

When you open the page, graphical representations of all the data centers and custom data centers in your environment appear. By default, they are shown in order of time remaining, beginning from the upper left, where the most constrained data centers appear. To identify possible oversized and undersized VMs in a data center, click its graphic. The area following refreshes to display details about the selected data center.

"Oversized VMs" displays the number of VMs determined to be oversized based on policies previously set. A chart details suggested reductions in the overall number of CPUs and GBs of memory and shows the percentage of total resources the reductions represent. Similarly, "Undersized VMs" indicates the number of VMs considered to be undersized, with a chart listing suggested increases in CPU and memory.

The table at the bottom of the page provides important information about the VMs. Table headings are Oversized VMs and Undersized VMs. VMs under each heading are grouped by cluster. Click the chevron to the left of a cluster name to list all the oversized or undersized VMs, respectively, in that cluster. You can check the box next to one or more VM names and click the **EXCLUDE VM(S)** button to prevent those VMs from being included in a resizing action. You can also select individual VMs to resize before clicking the **RESIZE VM(S)** button.

## Run a Rightsize Action on Oversized VMs

Run the action as follows:

1. Click the Oversized VM tabs.
2. Select the VMs you want to exclude from the action, if any and click **EXCLUDE VM(S)**. In the confirmation dialog box, click **EXCLUDE VM(S)**.
3. Select the VMs you want to include in the resizing action, or click the checkbox next to VM Name to include all VMs.
4. Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested reductions for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.

   **NOTE**
   Operational Actions must be activated in the vCenter cloud adapter instance.

5. Select the checkbox at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

## Run a Rightsize Action on Undersized VMs

Run the action as follows:

1. In the table headings, **Select** Undersized VMs.
2. Select the VMs you want to exclude from the action, if any and click **EXCLUDE VM(S)**. In the confirmation dialog box, click **EXCLUDE VM(S)**.
3. Select the VMs you want to include in the resizing action, or click the checkbox next to VM Name to include all VMs.
4. Click **RESIZE VM(S)**. The Resize VM(S) workspace appears. The table displays suggested increases for vCPU and memory. **Click** the edit icons to accomplish to changes you wish.
5. Select the checkbox at the bottom of the screen to indicate your understanding that, because workloads must restart to accommodate resizing, some work may be interrupted.

**Table 119: Rightsize Options**

| Option | Description |
|---|---|
| Select a data center. | Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object. |
| **ALL DATACENTERS \| X** | Toggle: click **ALL DATACENTERS** on the upper right when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers. |
| View: | Filter results to include data centers, custom data centers, or both. Option appears when you select **ALL DATACENTERS** on the upper right. |
| Group BY: | Filter results by criticality (least time remaining data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. Option appears when you select **ALL DATACENTERS** on the upper right. |
| Sort by: | Options (Options appear when you select **ALL DATACENTERS** on the upper right):<br>• Alarm clock graphic - list data centers/custom data centers by time remaining.<br>• Dollar sign - list data centers/custom data centers by potential cost savings.<br>• Scales graphic - list data centers/custom data centers by level of optimization. |
| **Select data center or ADD NEW CUSTOM DATACENTER.** | Options (Options appear when you select **ALL DATACENTERS** on the upper right):<br>• Select a data center from the carousel across the top of the page. All data refreshes with information for the selected object.<br>• Select **ADD NEW CUSTOM DATACENTER** to display a dialog box that allows you to define a custom data center. |
| Oversized VMs display | Displays the number of VMs identified as oversized, with suggested reductions for vCPU and memory size. |
| Undersized VMs display | Displays the number of VMs identified as undersized, with suggested increases for vCPU and memory size. |
| Table of Oversized and Undersized VMs | Tabular representation of the Oversized and Undersized VMs in the selected data center.<br><br>Click one of the headings - Oversized VMs or Undersized VMs - to refresh the table with data for that heading. The table lists the relevant VMs. To see the VMs hosted in a given cluster, click the chevron to the left of the cluster name.<br><br>Click the check box next to the VMs you want to act on, or click the check box next to the column heading VM Name to act on all the VMs.<br><br>Click the expand icon next to a VM to see an explanation of how VMware Cloud Foundation Operations computes the recommended size. The visual representation shows you the historical utilization pattern, the projection and the recommendation for rightsizing. You can change the graphical representation to show the history from one month to 12 months. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>The rightsizing recommendation is influenced by the risk level that you set in the policy. You can modify the risk level in the policy to change the recommendation.<br><br>Once you select a VM or VMs, the dimmed options above the table become visible, as follows.<br><br>**Exclude VM(s)**: the selected VMs are excluded from your subsequent action. Excluding VMs from a reclamation action can reduce the potential cost savings.<br><br>For Oversized VMs:<br>• **SCHEDULE ACTION**: Displays a dialog box enabling you to schedule one or more resize actions for oversized VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>• **RESIZE VM(s)**: The system displays a dialog box with suggestions for reducing vCPUs and memory. Click the edit icons to change resource size.<br>• **EXCLUDE VM(s)**: Excludes the selected VMs.<br>• **EXPORT ALL**: Exports the list of powered off VMs into a CSV file.<br><br>For Undersized VMs:<br>• **SCHEDULE ACTION**: Displays a dialog box enabling you to resize VM actions for undersized VMs. Expand the cluster name displayed in the table and select one or more VMs. Then, from the SCHEDULE ACTION drop down menu, select an action to be performed later. In the dialog box, you configure the schedule for the job. Scheduled jobs can be managed in Automation Central.<br>• **RESIZE VM(s)**: The system displays a dialog box with suggestions for increasing vCPUs and memory. Click the edit icons to change resource size.<br>• **EXCLUDE VM(s)**: Excludes the selected VMs.<br>• **EXPORT ALL**: Exports the list of powered off VMs into a CSV file.<br><br>**SHOW\|HIDE EXCLUDED VMS**: toggle displays or hides the list of VMs you previously excluded.<br><br>**INCLUDE VM(s)**: include the selected VMs in the actionable list. |

### Using Workload Optimization to Improve Performance

Workload Optimization in VMware Cloud Foundation Operations focuses on intelligently balancing and optimizing workloads across your virtualized infrastructure. The goal is to ensure that resources are distributed efficiently to maintain optimal performance, prevent resource contention, and make the best use of available hardware resources.

Using Workload Optimization, you can rebalance virtual machines and storage across clusters, relieving demand on an overloaded individual cluster and maintaining or improving cluster performance. You can also set your automated rebalancing policies to emphasize VM consolidation, which potentially frees up hosts and reduces resource demand.

Workload Optimization further allows you to potentially automate a significant portion of your data center compute and storage optimization efforts. With properly defined policies determining the threshold at which resource contention automatically runs an action, a data center performs at optimum.

Starting with version 8.6, you can run workload optimization on a custom data center which has clusters aross multiple data centers within a single vCenter instance. The prequisite for this is that the hosts in the different clusters must be under the same network. This means that the port groups must be the same across the data centers. You can activate cluster level optimization across data center boundaries by enabling the setting in the business intent workspace.

**VMware Aria Automation Integration**

When you add an instance to a VMware Aria Automation adapter or solution pack as well as to a Key definition for "vCenter_server" not found in the DITA map. adapter instance that is connected to the VMware Aria Automation server, using VMware Aria Automation-managed resources, VMware Aria OperationsVMware Cloud Foundation Operations automatically adds a custom data center for the Key definition for "vCenter_server" not found in the DITA map., using VMware Aria Automation-managed resources.

On the VMware Aria OperationsVMware Cloud Foundation Operations side, to get the day2 chain configured, you must make the following initial configurations:

1. In vCenter Server, **Administration -> Solutions** and then add the VMware vSphere adapter instance for the vCenter Server that is configured as an endpoint in VMware Aria Automation Server.
2. In vCenter Server, **Administration -> Solutions** and then add the VMware Aria Automation adapter instance for the server that will appear in the VMware Aria OperationsVMware Cloud Foundation Operations and VMware Aria Automation integration day2 chain.

VMware Aria OperationsVMware Cloud Foundation Operations can manage workload placement and optimization for the custom data centers that reside in VMware Aria Automation-managed clusters.

However, VMware Aria OperationsVMware Cloud Foundation Operations is not permitted to set tag policies for the custom data center. (At the Workload Optimization screen, the Business Intent window is not operational for VMware Aria Automation custom data centers.) When rebalancing a VMware Aria Automation custom data center, VMware Aria OperationsVMware Cloud Foundation Operations uses all applicable policies and placement principles from both systems: VMware Aria Automation and VMware Aria OperationsVMware Cloud Foundation Operations. For complete information on creating and managing VMware Aria Automation custom data centers that are managed by VMware Aria OperationsVMware Cloud Foundation Operations, see the VMware Aria Automation documentation.

### Configuring Workload Optimization

Workload Optimization offers you the potential to automate fully a significant portion of your cluster workload rebalancing tasks. The tasks to accomplish workload automation are as follows:

1. Configure the Workload Automation Details. See Workload Automation Details.
2. If you do not use the AUTOMATE function in the Optimization Recommendation pane at the Workload Automation screen, configure the two Workload Optimization alerts to be triggered when cluster CPU/memory limits are breached, and configure them as automated. When the alerts are automated, the actions calculated by Workload Optimization are run automatically. See Configuring Workload Optimization Alerts

**Prerequisites**

Workload Optimization acts on objects associated with the VMware vSphere Solution that connects VMware Aria OperationsVMware Cloud Foundation Operations to one or more vCenter Server instances. The virtual objects in this environment include a vCenter Server, data centers and custom data centers, cluster compute and storage resources, host systems, and virtual machines. Specific requirements:

- Workload Optimization is supported when the optimization candidate clusters are in the same network or port group.
- Optimization candidate clusters should have one of the following datastore configurations:
  - Shared Datastore
  - Datastore Cluster
  - vSAN Datastore
- A vCenter Adapter configured with the actions activated for each vCenter Server instance.
- A vCenter Server instance with at least two datastore clusters. sDRS must be activated and fully automated for datastore clusters configuration only.
- Any non-datastore clusters must have DRS activated and fully automated
- Storage vMotion must be set to ON at Workload Automation Details. The default is On.
- You must have permission to access all objects in the environment.

**Design Considerations**

The following rules constrain the possible computer and storage resource moves that can be performed.

> **NOTE**
> When VMware Aria OperationsVMware Cloud Foundation Operations suggests that you optimize clusters in a data center, the system does not guarantee it can run an optimization action. VMware Aria OperationsVMware Cloud Foundation Operations analytics can determine that optimization is desirable and can create a rebalancing plan. However, the system cannot automatically identify all the architectural constraints that may be present. Such constraints may prevent an optimization action, or cause an action in progress to fail.

- Moving compute and storage resources is allowed only within, not across data centers or custom data centers.
- Storage resources cannot be moved across non-datastore clusters.
- Compute-resource-only moves are permitted through shared storage.
- Virtual machines defined with affinity rules or anti-affinity rules are not to be moved.
- Virtual machines cannot be moved when residing on a local datastore, unless a storage swap exists on the local datastore.
- Virtual machines cannot be moved if they have data residing across multiple datastore clusters. Compute-only moves with similar shared storage are not permitted.
- A virtual machine cannot have data that resides across different storage types. For example, if a virtual machine has a VM disk on a datastore and a second VM disk on a datastore cluster, the virtual machine does not move, even when the datastore is shared with the destination or has swap on it.
- A virtual machine can use RDM so long as the destination datastore cluster can access the RDM LUN.
- A virtual machine can implement VM disks on multiple datastores inside a single datastore cluster.
- Workload Optimization may suggest moving virtual machines that are protected by vSphere Replication or Array Based Replication. You must ensure that all the clusters within a selected data center or custom data center have replication available. You can set up DRS affinity rules on virtual machines that you do not want moving across clusters.

### Workload Placement Page

Workload Placement lets you to optimize virtual machines and storage across clusters to reduce resource contention and maintain optimum system performance.

### Where You Find Workload Placement

Select **Workload Placement** under **Optimize** in the left menu.

### Workload Placement Page Options

In the Workload Optimization page, you see a list of data centers in a carousel, listed under three categories:

- Critical
- Normal
- Unknown

After you select a data center, you see the **ALL DATACENTERS** button on the upper right. Click **ALL DATACENTERS** when you want to switch the view to a filtered list of all data centers. Click **X** to return to a carousel view of data centers.

**Table 120: Workload Placement Page Options**

| Option | Description |
|---|---|
| View: | Filter results to include data centers, custom data centers, vRA-managed custom data centers, or all three. (Option appears if you select **ALL DATACENTERS** on the upper right.) |
| Group By: | Filter results by criticality (most out of balance data centers/custom data centers listed first) or by the vCenter Server to which each data center belongs. (Option appears if you select **ALL DATACENTERS** on the upper right.) |
| Sort By: | Options (Options appear if you select **ALL DATACENTERS** on the upper right):<br>• Alarm clock graphic - list data centers/custom data centers by time remaining.<br>• Dollar sign - list data centers/custom data centers by potential cost savings with capacity optimization.<br>• Scales graphic - Optimized. |
| Select data center or **ADD NEW CUSTOM DATACENTER** | Options (Options appear if you select **ALL DATACENTERS** on the upper right):<br>• Select a data center from the carousel across the top of the page. All data following refreshes with information for the selected object.<br>• Select **ADD NEW CUSTOM DATACENTER** to display a screen that enables you to define a custom data center. |

### Data Center Options

After you select a data center from the carousel, you see the following information and options.

> **NOTE**
> If you point your cursor to the lower right of a data center graphic, a tooltip may appear to let you know that the data center is using automated optimization.

### Optimization Status Tab

Appears when you select a data center or custom data center from the top of the screen.

**Table 121: The Optimization Recommendation Card**

| Option | Description |
|---|---|
| Status | • Optimized - indicates that workloads are optimized based on the settings you entered in the neighboring Operational Intent window, with no tag violations based on the settings you entered in the Business Intent window.<br>• Not Optimized - indicates that one of the following conditions is true: workloads are not optimized based on the settings you entered in the neighboring Operational Intent window AND/OR there are tag violations based on the settings you entered in the Business Intent window. In the event of tag violations, the offending tags are listed. |
| OPTIMIZE NOW | Runs optimizing actions based on the settings you entered in your Operational and Business Intent settings. |
| SCHEDULE | Displays a dialog box enabling you to schedule one or more optimization actions. If schedules are currently set for data center or custom data center optimization, a check mark appears next to the data center or custom data center name. |
| AUTOMATE | Continually seeks optimizing opportunities for data center or custom data center, based on the settings in the neighboring Operational Intent window or Business Intent windows. Scheduled optimizations are turned off while automatic optimization is on. Also, automated alerts are not operational when automatic optimization is on. Once you confirm automation, the system displays message, for example, 1) "Workload Optimization is looking for opportunities to automate," 2) "Your workloads are optimized according to your settings." or 3) "No eligible moves were found within the max number of compatibility checks allowed."<br>**NOTE**<br>To initiate Automation, you must have privileges for Environment -> Action -> Schedule Optimize Container. |
| TURN OFF AUTOMATION | Stops automatic optimization. Any scheduled optimizations come back online. |

**NOTE**
Sometimes an optimizing action may be recommended, for example to consolidate two hosts, but when you run the optimization, the generated placement plan does not show any potential consolidation. The seeming inconsistency results from the fact that recommended optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation can incur stress in the future, then consolidation is not recommended.

**NOTE**
Scheduled jobs created from the Workload Placement page will be shown in the Automation Central page. However, these jobs cannot be edited, cloned or previewed, but can only be activated or deactivated.

**Table 122: The Operational Intent Card**

| Option | Description |
|---|---|
| Utilization Objective | Indicates the main attribute of your current automation policy settings. Values are moderate, consolidate, or balance. |
| Edit | Displays the Workload Automation Policy Settings, where you can adjust settings for optimization and cluster headroom. For more information, see the topic, Workload Automation Details. |

**Table 123: The Business Intent Card**

| Option | Description |
|---|---|
| Intent | Allows you to define zones of infrastructure within cluster boundaries. |
| Edit | Displays a workspace where you can select criteria for placement of VMs. |

**Table 124: Details for Are your clusters meeting your utilization objective?**

| Option | Description |
|---|---|
| Are your clusters meeting your utilization objective? | Displays a table which presents data in the following columns based on the Demand model or Allocation model:<br>• Name<br>• CPU Workload<br>• Memory Workload<br>• DRS Settings<br>• Migration Threshold<br>• Violated Tags<br>• VM Name<br><br>To activate the Allocation model, set the overcommit ratios in the active policy. For more information, see the topic, Allocation and Demand Model in Workload Optimization.<br><br>Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster. The violated tags shows which clusters or host groups are breaching the business intent. The VM Name column shows the name of the VMs and tag value due to which tag violation is happening. Provides the option to set the DRS automation level for individual objects. |
| VIEW DRS SUMMARY | Select a cluster in the list, then click this link to display a page containing metrics for DRS performance and cluster balance in the selected data center. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| SET DRS AUTOMATION | Select a cluster in the list, then click this link to set the level of the DRS automation for the cluster. Note that clusters must be fully automated in order for workload optimization alerts to run actions set in the policies. |

**History Tab**

Displays a graphical depiction of executed manual and automated optimizations for clusters in the selected data center or custom data center, based on parameters you provide.

**Table 125: Details for History**

| Option | Description |
|---|---|
| Selected WLP process drop-down | The optimization action whose details you want to display. |
| Time duration drop-down | Last *n* hours - select the time parameter: last 6, 12, 24 hours or last 7 days. |
| Quick filter | choose a cluster name to search on. |
| Squares graphic | toggle between viewing processes in icon or circle form. |
| Circle | Toggle between viewing processes presented in a circle or on a straight line. |
| Back arrow - reset action. | Reset action. |

If you point your cursor to a specific cluster as displayed on the screen, the details of the cluster appear in a tool tip. Click the note card icon on the lower right of the tool tip to go to the Details screen for the cluster. When displayed in the circle format, rings in the circle indicate how much CPU and how much memory was used at any given time. For example, if memory usage was higher than recommended based on your policy settings, the memory circle appears red.

Note the timeline across the bottom of the screen. When you choose parameters, for example, WLP process name, time parameter and cluster name, indicators appear along the timeline, showing when processes were initiated.

To zero in on a specific event, choose a process from the drop-down menu. You can also click points on the marker floating above the timeline, which causes a descriptive tool tip to appear, then double-click the 'Double-click to zoom' icon on the lower right.

If the event you choose includes an actual movement of VMs, you see a blue ball containing the number of VMs moved and showing the direction of the move and starting and ending clusters.

**Optimization Potential Tab**

When you run Workload Optimization, VMware Aria OperationsVMware Cloud Foundation Operations runs compatibility checks and excludes those VMs which have constraints, and only optimizes resources of those VMs which can be moved. If you want to see the total potential of your workload optimization, assuming that all VMs can be moved, click the **CALCULATE OPTIMIZATION POTENTIAL** button in the Optimization Potential tab. Optimization Potential disregards the underlying constraints and recommends moves before the compatibility checks. You can download the report to see more details. Storage DRS must be set to fully automated.

If you want to see what can be realistically optimized, click **OPTIMIZE NOW** in the **Operation Status** tab. After you click **OPTIMIZE NOW**, you can download a report to review incompatibilities.

The optimization potential report helps you understand the difference between the optimization achievable when you run **OPTIMIZE NOW** and the total optimization potential.
See also Example: Run Workload Optimization

**Configuring Workload Optimization Alerts**

VMware Aria OperationsVMware Cloud Foundation Operations provides two preconfigured alerts designed to work with the Workload Optimization feature. You must take additional action in the Policies area to turn on the alerts and automate them so that predetermined actions are run when the alerts fire.

Ensure that you have all required permissions to access the Workload Optimization UI pages and manage vCenter Server objects.

The following preconfigured alerts are designed to work with the Workload Optimization feature:
- Data center performance can potentially be optimized in one or more clusters.
- Custom data center performance can potentially be optimized in one or more clusters.

The preconfigured alerts fire only if the AUTOMATE function is not turned on at the Workload Optimization screen which can be accessed by clicking **Workload Placement** under **Capacity** in the left menu.

1. Select **Operations › Configurations › Policy Definition**.
2. Select the policy that includes settings for the relevant data centers and custom data centers, for example, **vSphere Solution's Default Policy**.
3. Click the **Vertical Ellipses** next to **Add** and then **Edit**.
4. Click the **Alerts and Symptoms** tile.
5. Search on "can potentially be optimized" to locate the two alerts you want.
6. The alerts are activated by default/inheritance (see the **State** column).
7. The alerts are not automated by default/inheritance (see the **Automate** column). To automate the alerts, click the menu symbol to the right of the inherited value and select the green check mark.

Workload Optimization is fully automated for your environment.

To confirm that actions are taken automatically, monitor rebalance activity at the Workload Optimization screen.

**Workload Automation Policy Settings**

Provides options for refining policy settings specifically for Workload Optimization.

**Where You Find Workload Automation Settings**

Access this screen through the Policies pages:

1. Select **Operations › Configurations**.
2. Click the **Policy Definition** tile.
3. Select a policy that you want to modify. Ideally, this should be an active policy. Or, click the **ADD** button to add a new policy.
4. Select the **Workload Automation** card to review the changes, or click **EDIT POLICY** to make changes.

,

Refer to Workload Automation Details.

## Business Intent Workspace

You can use vCenter Server tagging to tag VMs, hosts, and/or clusters with specific tags. VMware Aria OperationsVMware Cloud Foundation Operations can be configured to leverage tags to define business-related placement constraints. Using tags, you can place VMs on hosts/clusters with matching tags, or exclude VMs and clusters from participating in Workload Placement.

### Where You Find Business Intent

Open the **Workload Placement** page from under **Capacity** in the left menu. On the Workload Placement page, select a data center or custom data center from the top row, and click **Edit** in the Business Intent window.

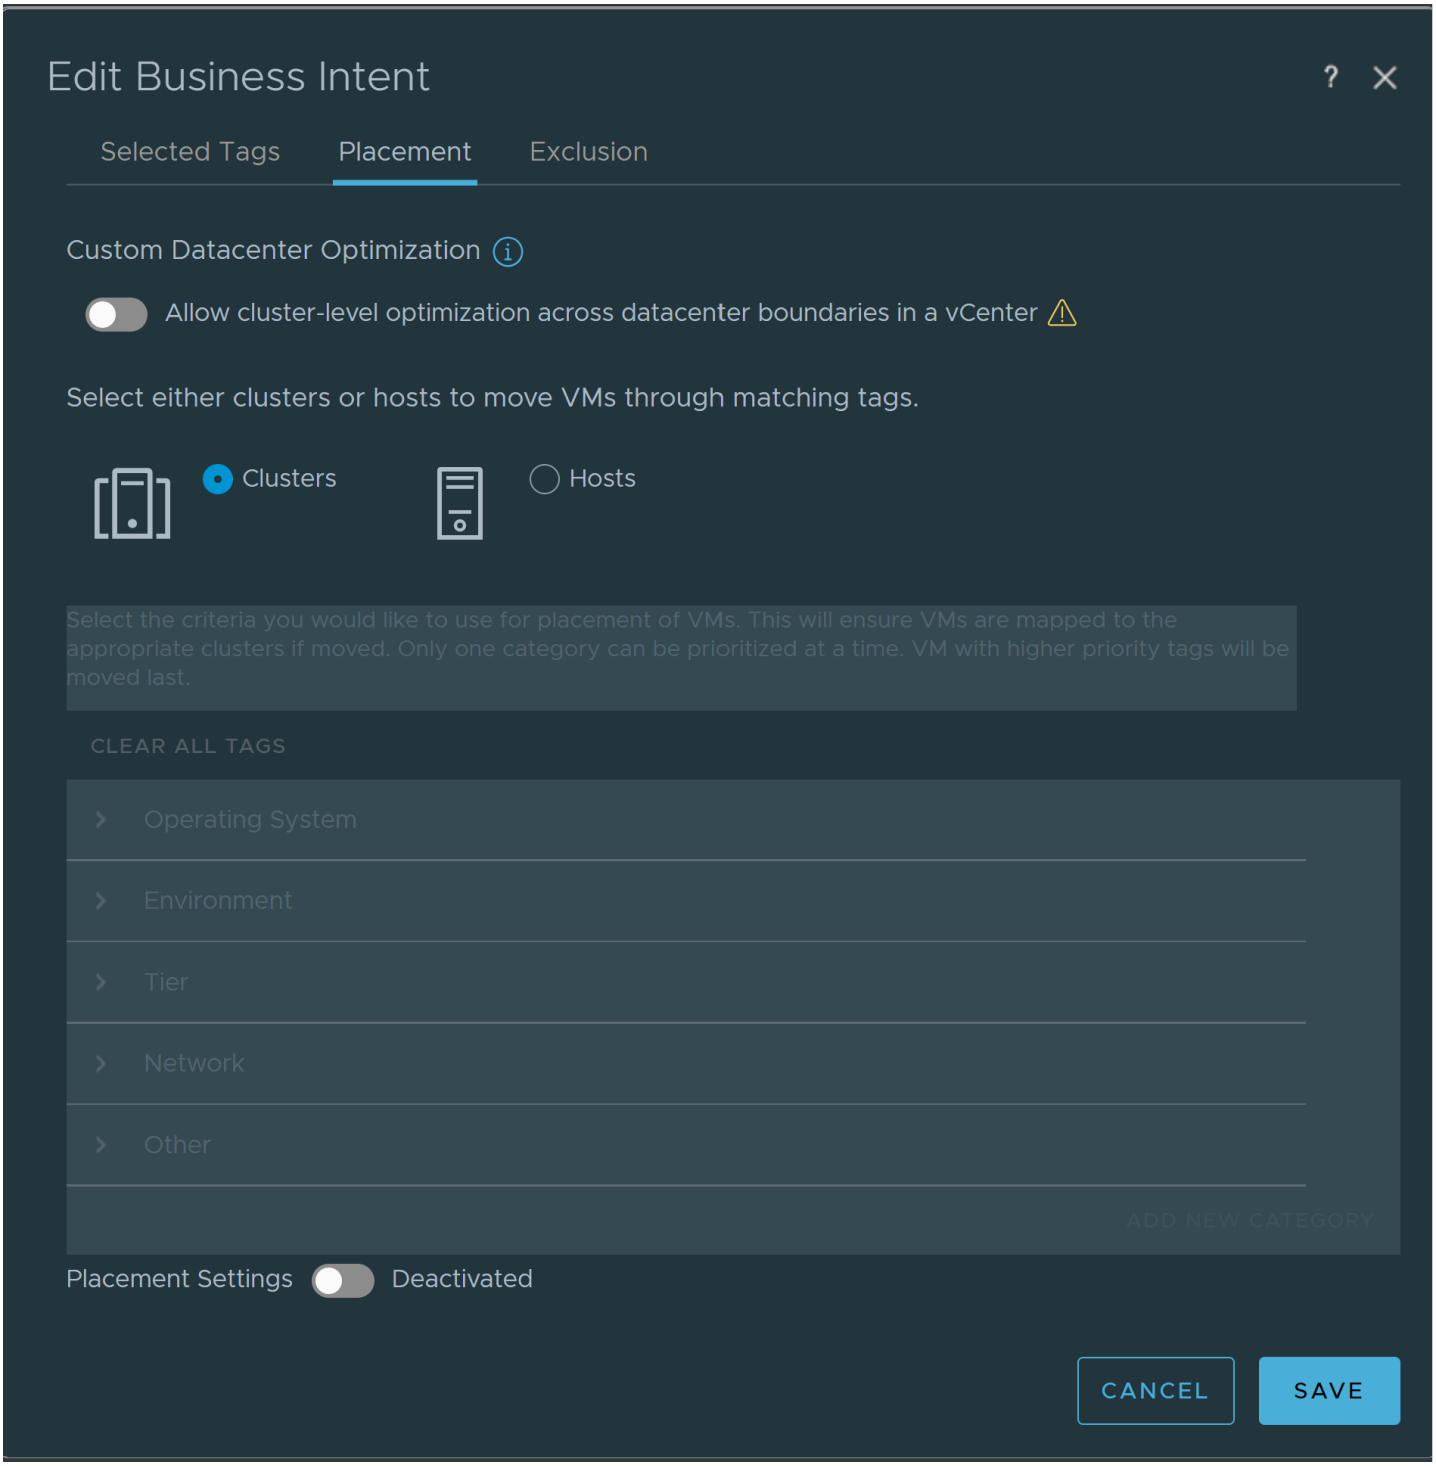


## How the Edit Business Intent UI Works

The Edit Business Intent dialog box has three tabs. The **Placement** tab is where you select the tags which VMware Cloud Foundation Operations uses for Workload Placement. The **Exclusion** tab is where you specify tags to exclude VMs and Clusters from participating in the Workload Placement. For example, you may have VMs running critical workloads in a large environment. These VMs may be running on specific hosts so that they always get the resources they need for running the workloads. You can exclude such VMs from Workload Placement if you do not want the VMs from moving

from their original location. The clusters which are excluded using the exclusion tags are not displayed in the Workload Placement page, and are not considered in the Workload Placement calculation.

**NOTE**
Tag exclusion was introduced newly in version 8.16.

The tags that you select in the **Placement** and **Exclusion** tabs are displayed in the **Selected Tags** tab.
You can activate cluster level optimization across data center boundaries in a vCenter Server if you have created a custom data center across multiple data centers within a single vCenter Server. You must ensure that the prerequisite networking requirements for movement of VMs are met before running workload optimization on such a custom data center.

To edit Business Intent values, you must have the necessary permissions. When you are logged in with administrative privileges, click Administration on the left menu and go to **Roles** under **Access Control** . Select the name of the role to which you want to provide permissions and then click **Edit** in the permissions section. Select the **Read** and **Write** checkbox under Administration → Configuration → WLP Settings.

## Establishing Business Intent

Tags are implemented in vCenter Server as *key:value* labels that allows operators to add meta-data to vCenter Server objects. In vCenter Server terminology, the *key* is the tag category and the *value* is the tag name. Using this construct, the tag OS: Linux can indicate a cluster or VM that is assigned to the category OS with a tag name of Linux. For complete information on vCenter Server tagging capabilities, refer to the vCenter Server and Host Management guide.

To specify tags that are either considered for placement, or excluded from placement, first go to the appropriate tab in the Edit Business Intent dialog box.

Select the radio button for the type of object you want to associate with VMs in this business intent session: Clusters or Hosts. Then, click the **Placement Settings** or **VM Exclusion Settings** toggle button.

The system provides several suggested categories. These categories are only suggestions. You must specify the actual categories in vCenter Server after you expand the section for a suggested category . For example, in section "Tier", you can specify the actual vCenter Server tag category that represents tier semantics, for instance, "service level".

- Operating System
- Environment
- Tier
- Network
- Other

Any actual categories you specify must first be created in vCenter Server.

Then you can associate tagged VMs with clusters or hosts, based on the rules for each type of tagging.

1. Click the chevron to the left of the first suggested category. A **tag category** field appears.
2. Click the drop-down menu indicator and choose a category from the list defined in vCenter Server.
3. Click the drop-down menu indicator in the Tag Name (Optional) field and choose a tag name from the list defined in vCenter Server.
4. Click **Include Tag**. All VMs with that tag are associated with the category.

## Rules for Host-Based Placement

To set host level placement constraints, VMware Aria OperationsVMware Cloud Foundation Operations automatically creates and manages DRS rules. All conflicting user-created DRS rules are deactivated.

These rules include the following:

- Any VM-VM affinity and anti-affinity rules.
- Any VM-Host affinity and anti-affinity rules.

You must check the selection box next to the statement, "I understand that VMware Cloud Foundation Operations will deactivate all my current and future DRS rules".

### View DRS Summary

The View DRS Summary page provides insight and perspective into the actions DRS is taking to balance a cluster. You can view DRS settings for the cluster and cluster balance metrics, and determine if recent vMotions are DRS- or user-initiated.

#### Where You Find the View DRS Summary Page

In the left menu, click **Workload Placement** under **Capacity**. Then select a cluster name in the "Are your clusters meeting your utilization objective?" sedction. The dimmed View DRS Summary and Set DRS Automation links turn live. Click the link to display the DRS summary information.

**Table 126: DRS Summary Values**

| Pane/fields | | Value |
|---|---|---|
| <*cluster name*> | | Name of the selected cluster |
| | Automation Level | Enabled/Disabled. DRS is running or not. |
| | Migration Threshold | Aggressive/Default/Moderate |
| | Active Memory Used | False/ *nn*% |
| Cluster Balance | | Shows the variations in the DRS cluster balance metric over time as DRS runs. The graph shows how DRS reacts to and clears any cluster imbalance each time it runs. |
| | Cluster Imbalance | The range of potential imbalance values, as expressed in vCenter DRS metrics. |
| | Total Imbalance | The level of imbalance in a cluster, as measured by vCenter DRS metrics. |
| | Tolerable Threshold | The upper limit of what is tolerable in cluster imbalance. Designated by a green dotted line, this is a vCenter DRS metric. |
| VM Happiness | | A bar graph summarizing the total happy and unhappy VMs in the cluster. For individual VMs, there is a presentation of performance metrics related to its happiness, such as %CPU ready time and memory swapped. |
| | Happy VMs | Total of happy VMs are shown in green. Click in the green zone to show a list of these VMs in the Happy/Unhappy VMs pane to the right. |
| | Unhappy VMs | Total of unhappy VMs is shown in red. To show a list of these VMs in the Happy/Unhappy VMs pane to the right, click in the red zone . |
| Happy/Unhappy VMs | | Lists by name all the VMs in the zone you clicked in the VM Happiness pane. |
| VM Metrics | | Shows the trend in VM happiness or unhappiness |
| Recent vMotions | | The number of recent vMotions, plotted against time. |
| vMotion Details | | Shows the number of DRS-initiated and user (non-DRS) initiated vMotions over time. You can choose which type you want to view. |
| | Date/VM | Date of a given vMotion. |

*Table continued on next page*

*Continued from previous page*

| Pane/fields | | Value |
|---|---|---|
| | Source/Destination | Source and destination of moved VMs. |
| | Type | DRS-initiated or user initiated. |

### Optimization Schedules

Use the Optimization Schedules page to edit or delete optimization schedules that you set up in the Manage Optimization Schedule Dialog Box at the Workload Optimization main screen.

#### Where You Find Optimization Schedules

- **Optimization Schedules** is available in the left menu, **Workload Placement** under **Capacity**.
- At the Workload Placement Page page, select in the data center whose optimization schedule you want to edit or delete. Then click **SCHEDULE** in the Optimization Recommendation pane.

**Table 127: Optimize Schedules Options**

| Option | Description |
|---|---|
| Edit icon | Select a schedule from the list, then click the **Edit** icon.<br><br>The Manage Optimization Schedule Dialog Box appears, with the data for the selected schedule filled in. |
| Delete icon | Select a schedule from the list, then click the **Delete** icon.<br><br>The selected schedule is deleted and does not run. |

See also Example: Run Workload Optimization

### Manage Optimization Schedules

Allows you to set up a regular schedule for optimizing a selected container.

#### Where You Find Manage Optimization Schedules

At the Workload Optimization screen, select **SCHEDULE** from the pane: Optimization Recommendation

| Option | Description |
|---|---|
| Schedule Name | Meaningful name for the schedule |
| Time Zone | Choose the time zone for the action |
| Recurrence | Indicate how often you want the optimize action to run. Complex schedules can be defined, for example, select the Monthly option and choose to run the action on Tuesdays and every other Thursday, beginning on the fifth of the month. |
| Start on: | Day to start the optimization schedule. |
| Start at: | Time to start the optimization schedule. |

*Table continued on next page*

*Continued from previous page*

| | |
|---|---|
| Expire after: | Designate a set number of scheduled runs. |
| Expire on: | Designate an exact date for the actions to end. |

See also Example: Schedule a Repeating Optimization Action

**Optimize Placement**

A two-page dialog box that provides information about optimizing the workload of a selected container. When you run the optimization action, VMware Aria OperationsVMware Cloud Foundation Operations checks which of the VMs can be moved to a different cluster for better optimization of resources, based on the settings you entered in your Operational and Business Intent settings. You can download a report that provides information about the list of VMs that were included in, and excluded from, the move plan. The report provides reasons as to why some VMs were excluded from the plan.

First page: The current workload ("before," for example, CPU 105%) and projected results ("after," for example storage utilization 45%) for a possible optimizing action.

Second page: The exact moves planned for compute and storage resources.

> **NOTE**
> It is possible that there is no optimization move plan. Review the report to see why VMware Aria OperationsVMware Cloud Foundation Operations could not provide a move plan.

**Where You Find Optimize Placement**

At the Workload Optimization screen, select OPTIMIZE NOW in the Optimization Recommendation pane.

**Table 128: Optimize Clusters Options**

| Option | Description |
|---|---|
| Compare Cluster Balance | If you are satisfied with the before and after numbers (First page, above), click NEXT. |
| Review Optimization Moves | If you are satisfied with the moves planned (Second page, above), click BEGIN ACTION.<br>**NOTE**<br>Review the optimization plan report before you click BEGIN ACTION. |
| Download Report | The optimization plan report is in CSV format, and provides the following information:<br>• Summary of the optimization plan.<br>• Summary of the moves that make up the optimization plan.<br>• Issues related to the data center. Resolve these issues before proceeding with the optimization.<br>• Issues and incompatibilities applicable to specific VMs and their configurations. Resolve these issues, if applicable.<br>• Failed move attempts applicable to the specific VMs and their target destinations, as determined from the VM move plan. Resolve these issues and incompatibilities. |

See also Example: Run Workload Optimization.

## Using Workload Optimization

Use the Workload Optimization UI pages to monitor optimizing moves in a fully automated system. If your system is not fully automated, you can use the UI to conduct research and run actions directly.

VMware Aria OperationsVMware Cloud Foundation Operations monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization screen. Depending on what appears on the screen, you might use optimization functions to distribute a workload differently in a data center or custom data center. Or you may decide to perform more research, including checking the Alerts page to determine if any alerts have been generated for objects of interest.

For comprehensive general instructions on responding to alerts and analyzing problems related to objects in your environment, see Monitoring Objects in Your Managed Environment by Using vRealize Operations Manager.

For comprehensive general instructions on responding to alerts and analyzing problems related to objects in your environment, see the *VMware Aria OperationsVMware Cloud Foundation Operations User Guide*.

The following examples demonstrate the primary ways you can use Workload Optimization to keep your data centers balanced and performing their best.

### Example: Run Workload Optimization

As a virtual infrastructure administrator or other IT professional, you use Workload Optimization functions to identify points of resource contention or imbalance. In this example, you manually run an optimization action to consolidate demand.

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

1. When you log into VMware Aria OperationsVMware Cloud Foundation Operations, you see the Launchpad page. Click the **Capacity** card and then **View** in the **Workload Optimization** card.
   The Workload Optimization page appears. Data centers are grouped by Criticality, with the three troubled data centers appearing in a carousel across the top of the page: DC-Bangalore-18, DC-Bangalore-19, DC-Bangalore-20. A Not Optimized badge appears in the lower right corner of each graphic.
2. If no data center is preselected, select DC-Bangalore-18 from the carousel.

   Comprehensive data about the state of the data center follows.
3. Based on the available data, you determine an optimization action is required.

   CPU workloads can be consolidated such that a host in Cluster 3 can be freed up.

**Table 129: Panes and Widgets**

| Pane | Contents |
|---|---|
| Workload Optimization | Status shows as Not Optimized. A system message says, "You can consolidate workloads to maximize usage and potentially free up 1 host."<br>The message reflects that you have set policies to emphasize consolidation as a goal in optimization moves. The system is saying you can free up a host through consolidation. |
| Settings | The current policy is Consolidate. The system advises: Avoid Performance Issues, Consolidate Workloads. |

*Table continued on next page*

*Continued from previous page*

| Pane | Contents |
|---|---|
| Cluster Workloads | Cluster 1 CPU Workload is 16%. Cluster 2 CPU Workload is 29%. Cluster 3 CPU Workload is 14%. Cluster 4 CPU Workload is 22%. |

4. Click **OPTIMIZE NOW** in the Workload Optimization pane.
   The system creates an optimization plan, which depicts BEFORE and (projected) AFTER workload statistics for the optimization action.
5. If you are satisfied with the projected results of the optimization action, click **NEXT**.
   The dialog box updates to show the planned moves.
6. If you need more information about the VMs which are included or excluded in the plan, click **Download Report** to see the optimization plan. You can review the reasons for incompatibilities and why some VMs were excluded from the plan.
7. Optional: If you want to know the total optimization potential of the move, assuming that there were no incompatibilities and all your VMs can be included in the optimization plan, click **Cancel**, and go to the Optimization Potential tab in the Workload Optimization page. Click **Calculate Optimization Potential** to see the total optimization potential of your data center.
8. Review the optimization moves, then click **BEGIN ACTION**.
   The system runs the compute and storage resource moves.

The optimization action moved compute and storage resources from some clusters to other clusters in the data center, and so freed up a host on one cluster.

> **NOTE**
> The Workload Optimization page refreshes every five minutes. Depending on when you run an optimization action, the system might not reflect the result for up to five minutes, or longer when longer-running actions extend the processing time.

To confirm that your optimization action was completed, go to the Recent Tasks page by clicking **Administration** in the left pane. In the Recent Tasks page, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, first filter on Starting Time and scroll to the time when you began the action, then select the Object Name filter. Finally, enter the name of one of the VMs in the rebalance plan.

> **NOTE**
> Sometimes an optimizing action may be suggested, for example to consolidate two hosts, but when you run the optimization, the generated placement plan does not show any potential consolidation. The seeming inconsistency results from the fact that suggested optimization actions are based on current conditions, whereas the placement plan logic includes forecasting. If forecasting predicts that consolidation might incur stress in the future, then consolidation is not suggested.

### Example: Schedule a Repeating Optimization Action

As a virtual infrastructure administrator or other IT professional, you determine that compute and storage resources in a given data center are volatile and a regularly scheduled optimization action can address the problem.

Ensure that you have all required permissions to access the Workload Optimization UI and manage vCenter Server objects.

VMware Aria OperationsVMware Cloud Foundation Operations monitors virtual objects and collects and analyzes related data that is presented to you in graphical form at the Workload Optimization page. Depending on what appears, you may

determine that you must schedule optimization functions to distribute a workload more evenly in a data center or custom data center.

1. Select **Workload Placement** under **Capacity** in the left menu.
2. From the carousel of data centers across the top of the page, select a data center for which you want to schedule repeated optimization actions.
3. In the Workload Optimization pane, click **SCHEDULE**.
4. Give the schedule a name and choose a time zone.
5. Determine how often you want to repeat the optimization action and click the relevant **radio button** under Recurrence.
   Depending on your selection under Recurrence, additional options appear to the right. In this instance, you choose to repeat the optimization daily.
6. Leave the current date and time.
7. Select the **Repeat every day** radio button.
8. Select the **Expire after** radio button and tick the counter up to 6.
9. Click **Save**.

The optimization action repeats for six days, then stops.

At the Workload Optimization page, the Scheduled button appears in the upper right of the Workload Optimization pane if optimization actions are scheduled for the selected data center. If you want to edit or delete a schedule, click the **Scheduled** button. The Optimization Schedules page appears, where you can perform those actions.

> **NOTE**
> If you schedule a number of optimization actions close together, and the optimization plans of two or more actions include overlapping functions, that is, they impact the same set of resources, the system shifts the actions into a queue. As a result, some actions may complete later than expected, with longer running actions and other potential system constraints extending the lag time. Optimization actions that do not overlap can run concurrently.

To confirm that your optimization action was finished, go to the Recent Tasks screen. In the Recent Tasks screen, use the Status function on the menu bar to locate your action by its status. You can also search using a range of filters. For example, filter on Event Source and enter the name of the scheduled optimization plan.

> **NOTE**
> Because real-time data center resource contention is dynamic, the system calculates a new optimization plan each time the scheduled optimization action starts, but before it runs. The system does not run the action if the system determines that the data center container is balanced at this moment. On the Recent Tasks page, the name of the affected data center appears in the Object Name column, and the Message “The optimization of the selected container cannot be improved” appears under Details. Another possibility is that a scheduled optimization plan is attempted, but does not go forward. In this event - which is not the same as a "failed" action - the name of the affected data center also appears in the Object Name column.

## How to Plan for Capacity Changes

Planning for capacity changes in VMware Cloud Foundation Operations involves simulating hypothetical changes to the environment, such as adding new VMs or hosts, to understand their potential impact on resource utilization. This helps in proactive planning.

## What-If Analysis: Modelling Workload, Capacity, or Migration Planning

Using the What-If tool, you can plan for an increase or decrease in workload or capacity requirements in your virtual infrastructure. To evaluate the demand and supply for capacity on your resources, and to assess the potential risk to your current capacity, you can create scenarios for adding and removing workloads. You can also determine how much capacity you require to make a migration work. You can run one scenario or group scenarios and run them cumulatively.

### Why Create a Scenario

A scenario is a detailed estimation of the resources you must have available in your environment to incorporate upcoming changes. You define scenarios that can potentially add resources to actual data centers. VMware Aria OperationsVMware Cloud Foundation Operations models the scenario and calculates whether your desired workload can fit in the targeted data center. You can save multiple scenarios for comparison or review.

### Committed Scenarios

When you are sure that you need to reserve capacity, you can commit the scenario to have VMware Aria OperationsVMware Cloud Foundation Operations set aside resources for new, upcoming, or planned workloads. A committed scenario is a supposition about how the capacity and load change on your objects when you change the conditions in your virtual infrastructure environment. You do not have to implement the changes that your committed scenario represents. By committing a scenario, you can determine your capacity requirements before you implement the actual changes.

#### Why Create a Committed Scenario

In organizations which have separate capacity management and operations teams, committing a scenario helps stakeholders understand the current capacity and upcoming capacity requirements across the board. With committed scenarios, capacity is reserved and this prevents the operations team from performing adhoc resource increase on workloads, while the capacity manager is engaged in resource planning of new projects.

Committed Scenario also helps the team responsible for infrastructure expansion, as it provides actionable insights into future scenarios. In the event capacity becomes limited, it could be accounted for in the expansion.

### Where You Find What-If Analysis and Committed Scenarios

#### What-If Analysis

Click **What-If Analysis** in the left menu under **Capacity**and in the What-If Analysis page, you see a list of seven What-If Analysis panes in the **New** tab.

#### Committed Scenarios

In the left menu, click **Committed Scenarios** under **Capacity**. In the Committed Scenarios page, click the **ADD** button to directly create a committed scenario without first creating a What-If scenario.

The overview tab of the What-If Analysis page has seven panes. Each pane lets you run What-If scenarios to optimize capacity based on the following areas:

- Workload Planning: Traditional
- Workload Planning: Hyperconverged
- Infrastructure Planning: Traditional
- Infrastructure Planning: Hyperconverged
- Migration Planning: VMware Cloud
- Migration Planning: Public Cloud
- Datacenter Comparison: Private Cloud

## How What-If Analysis and Committed Scenarios Work

You can run What-If scenarios to see how much capacity will remain after you add or remove VMs or hosts and add hyperconverged infrastructure (HCI) nodes. Migration planning shows you the capacity and cost information after migrating to cloud based infrastructure.

Scenarios that you save for later are displayed as a list when you browse to the **What-if Analysis** and **Committed Scenarios** pages. You can run, edit or delete the saved scenarios. When you run a scenario, whose start date occurs in the past, you get a dialog box asking you if you would like to run the scenario with the current date. If you choose **No**, the scenario is not run. If you choose **Yes**, VMware Aria Operations runs the scenario with the current date as the start date. The end date is not affected.

Use the advanced filter and search box to look for saved scenarios by scenario name, scenario type, datacenter, and cluster. You can select more than one compatible scenarios and run them together. For example, you can create a scenario to remove hosts using the **Physical Infrastructure Planning** pane, because your organization has hardware that will soon become obsolete. You can create another scenario to add hosts to your physical infrastructure to account for new hardware that will replace the obsolete ones. You can run both these scenarios together to see the capacity after removing old hardware and adding new hardware.

You can only combine scenarios that pertain to the same object. Use the filters in the **Assess** page under **Capacity**to narrow down the list based on scenario name, type, data center, or cluster.

You can select the following combinations of scenarios and run them together:

**Workload Planning and Physical Infrastructure Planning**

- Add VMs
- Remove VMs
- Add Hosts
- Remove Hosts

## The Results Page

The Results page displays the results of running one or more saved scenarios. For Workload and Infrastructure Planning What-If scenarios, the summary page displays the allocation and demand values. If you do not see the allocation values, make sure that the overcommit ratios are activated in the policy. For more information, see the Policy Allocation Model Element topic in the Configuring guide.

To add or remove saved scenarios and run them again cumulatively, click **Edit** in the **Assess** page. To commit a scenario and reserve capacity, click the **COMMIT SCENARIO** button. The Create Committed Scenario fly-out opens from the right hand side of the page. Add a name to the scenario you want to commit. Provide and implementation date, and optionally, an end date and click **SAVE**.

### What-If Analysis - Workload Planning: Traditional

You define scenarios that can potentially add workloads to actual data centers. VMware Aria OperationsVMware Cloud Foundation Operations models the scenario and calculates whether your desired workload can fit in the targeted data center or custom data center. You can also define scenarios that can potentially remove workloads from data centers. VMware Aria OperationsVMware Cloud Foundation Operations calculates the time remaining and capacity remaining on the cluster when workloads are removed from the cluster.

## Where You Find What-If Analysis - Workload Planning: Traditional

Click **What-If Analysis** in the left menu under **Capacity**. Click **Add VMs** or **Remove VMs** in the Workload Planning: Traditional tile.

## How What-If Analysis - Workload Planning: Traditional Works

Capacity Optimization allows you to forecast successfully the impact of adding a workload to an application. By trying various scenarios, you can arrive at an optimum configuration. When you add VMs in the Workload Planning: Traditional pane, you can select the exact data center or custom data center where you want to locate the new workload. You can even pick a specific cluster where the workload is to reside.

In selecting the profile of your workload, you have two options:

- Configure the workload manually by specifying vCPUs, memory, storage, and expected use percentage. You have the further option to click Advanced Configuration and specify more precise characteristics for your workload.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system lets you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for the new workload, enter the start and end date for the period when you want the workload to be active. The default is: starting today and ending one year from today. The system can project scenarios ending up to one year from the current date.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the VMware Aria OperationsVMware Cloud Foundation Operations analysis and assessment of your plan.

The system lets you know immediately if the proposed workload fits or does not fit in the suggested location. If it fits, the results list the prime target cluster and any additional possible locations. The system also projects time remaining before the workload runs out of resources. If you select scenario details, the system displays a graphic depiction of resource use. For each attribute value - vCPU, memory, and storage - the amount by which the workload increases the percentage of total application capacity used is shown against a time line. The graph shows the existing percentage used in blue and the total of existing usage and added usage as a percentage of total capacity in green.

If the proposed workload does not fit, the system announces the outcome and provides the following information:

- How much the added workload reduces the time remaining for the target cluster, for example, from one year to zero.
- The discrepancy between the space available in the target cluster and what the proposed workload requires, for example, 100 GB of memory.
- The cost of the workload on the VMware Hybrid Cloud and on the public cloud.

## About Clouds

When you run a scenario in What-If Analysis, you get a recommendation based on cost relative to workload placement on different clouds. This cost-based recommendation varies for different clouds.

Private Cloud and VMware Cloud on AWS costs are computed based on resource usage levels.

Public clouds, AWS, IBM Cloud, Google Cloud, Microsoft Azure, and user-defined cloud costs are dependent on the selected configuration, that is, for the allocated resources. These public cloud instances are selected based on the close proximity rule, with simulated resource allocation values and in some scenarios, the exact configuration match available in the cloud instance list is not available. Due to this issue, these public cloud costs can be inherently higher in comparison.

## How What-If Analysis - Remove Workload Works

This feature of Capacity Optimization lets you to forecast successfully the impact of removing a workload. By trying various scenarios, you can arrive at an optimum configuration. Once you select the Workload Planning screen, you can select VMs from the concrete cluster data center or from the customer data center from which you want to remove the existing workload.

While removing workloads, you have two options to define the workload:

- Select existing VMs and use their projected utilization to evaluate the impact of removing workloads.
- Configure the workload manually by specifying the vCPUs, memory, storage, and expected use percentage.

Enter the start and end date for the period during which you want the workload to be removed. By default, the start date is today and the end date is one year from today. The end date is left empty by default. The system can project scenarios ending up to one year from the current date.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the VMware Aria OperationsVMware Cloud Foundation Operations analysis and assessment of your plan.

**Table 130: What-If Analysis Workload Page Options**

| Option | Description |
|---|---|
| Add/Remove VMs | Click **Add VMs** or **Remove VMs** to create a scenario for adding or removing workload. When clicked, the command displays the Add Workload or Remove Workload screen. |
| Scenario Name | In the heading of the Saved Scenarios table. Selecting the check box next to the name selects all scenarios in the list and turns on the dimmed **Delete** button. |
| Scenario type | Name of the scenario type. Values are Add Workload, Remove Workload, Add Capacity, Remove Capacity, and Migrate. |
| < *scenario_name*> | Name of a saved scenario. Selecting the check box next to a name turns on the dimmed **Run Scenario**,**Edit**, and **Delete** buttons. |
| All Filters | Use the filter to search for a specific scenario by name or type. |
| Show Columns | Click the small button on the lower left to display the Show Columns dialog box. You can select up to four columns to display in the table: Scenario Name, Scenario Type, Date Created, and Scenario Start and End Date. |

### Add or Remove VMs: Traditional

As part of the What-If workload planning for traditional infrastructure, Workload Planning: Traditional is the pane you use to fill in the details of your virtual machines. You select where to add or remove the workload, configure it yourself or use an existing VM as a template, and establish a time frame. You also have an advanced configuration option that lets you define your configuration more precisely.

#### Where You Can Add or Remove VMS

At the What-If Analysis screen, click **Add VMS** or **Remove VMS** in the Workload Planning: Traditional pane.

**Table 131: Workload Planning: Traditional Add VMs Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario |
| **Location** | Where do you want to add the workload? Select from the list of existing data centers. You can optionally select the exact cluster where you want the workload to reside. |
| **ApplicationProfile**/**Configure** | Allows you to configure the virtual compute resource, including vCPU, memory, and storage. |
| **Application Profile**/ **Import from existing VM** | Displays the Select VMs dialog box where you can select one or more existing VMs to use as templates for your workload. Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br><br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br><br>Click **OK** when done. |
| Choose Your Workload:<br>• **CPU**<br>• **Memory**<br>• **Disk space** | With the Configure radio button selected, you can size your workload by defining values for vCPU, memory, and disk space. These are your allocation values. |
| **Expected Utilization** | Set the projected percentage of total workload capacity you expect to average. Click **Advanced Configuration** to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning. VMware Aria OperationsVMware Cloud Foundation Operations calculates demand based on these values. |
| **Annual Projected Growth** | Set the percentage by which you expect your capacity go grow, annually. Click **Advanced Configuration** to set the percentage growth of CPU, Memory, and Disk individually.<br><br>For example, if the utilization is 100 at the start date, and you set the annual growth % to 10%, then at the end of the year the utilization will grow to 110.<br><br>The Annual Projected Growth can be set to 0% if no growth is expected. |
| Number of VMs (optional)/**Quantity** | You can optionally select how many VMs to spread the workload across. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. |
| **Run Scenario** | Click to run the scenario. The system calculates whether it fits into the location you selected. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

**Table 132: Workload Planning: Traditional Remove VMs Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario. |
| **Location** | From where do you want to remove the workload? Select from the list of existing data centers. You can optionally choose the exact cluster from where you want to remove the workload. |
| **ApplicationProfile**/**Configure** | Allows you to configure the virtual compute resource, including vCPU, memory, and storage.<br>After you have configured the scenario, enter the quantity of custom VMs that you want to remove. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| **Application Profile/Import Existing VMs** | Displays the Select VMs dialog box where you can choose one or more existing VMs. Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to remove from your workload.<br>**NOTE**<br>The recommended limit is 100 VMs as a maximum for workload removal. |
| **Application Profile** / **Custom**: Choose your workload<br>• **CPU**<br>• **Memory**<br>• **Disk space** | With the Configure radio button selected, you can size your workload by defining values for vCPU, memory, and disk space. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. You can also leave the end date blank. |
| **Run Scenario** | Click to run the scenario. The system calculates the impact on the cluster (time remaining and capacity remaining) when removing the workload. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

**Select VMs**

Use the **Select VMs** dialog box to choose the VMs whose attributes you want to copy or remove for your Workload Planning: Traditional or Workload Planning: Hyperconverged what-if scenarios.

**Where You Find Select VMs**

From the What-If Analysis screen, click **Add VMS** or **Remove VMS** in the Workload Planning: Traditional or Workload Planning: Hyperconverged pane. When you have entered a **Scenario Name** and **Location**, click the **Import from existing VM/Existing VMs** radio button, then click **Select VMS**. On the left is a selection box that allows you optionally to choose all VMs. To add a VM to the selected list on the right, double-click on the VM name. Following are the rest of your options:

**Select VMs**

| Option | Description |
|---|---|
| All Filters | Use the advanced filter and search option to look for VMs. Click the drop-down icon to narrow your search based on the available options. |
| Select (`nn`). | Select the VMs listed on the current page, from which to import, or remove characteristics. |
| Select all (`nn`) VMS | Click to select all the VMs across all the pages, based on the filters you have set. The number of VMs that you can select by clicking this option is limited to 500 VMs. |
| Selected | List of VMs you selected from RESULTS. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| OK | When you have selected the VMs you want, click **OK** to return to the Add Workload or Remove Workload screen, where your selected VMs are listed. |

Under Application Profile, in the Selected VMs table, enter the number of copies of each VM you selected to add or remove in the Quantity column.

### Scenario Summary Page for Workload Planning - Traditional

The scenario summary page for Workload Planning - Traditional displays the workload values after you add or remove VMs in a data center and cluster you specify.

#### Summary Section

The summary section displays the inputs that you provided for the Add VMS or Remove VMS scenario. You can change the workload location by selecting a different vCenter and cluster combination. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

#### Private Cloud and Public Cloud sections for Add VMS

The Private Cloud: Datacenter, Public Cloud, and Private Cloud sections display information which help you understand where your workload would fit, the associated costs, and the time remaining based on peak CPU, Memory, and Disk Space for demand and allocation model after you add VMs. You can choose to commit the Private Cloud: Datacenter results by clicking the **COMMIT SCENARIO** button. To see more details, click **VIEW DETAILS**.

You can see the cost per month for different private and public clouds. Click **Learn More** for a more detailed comparison based on your workload.

#### Scenario Results Section for Remove VMS

This section displays a graphical view of the projected CPU, Memory, and Disk Space values for the demand and allocation models after the VMs are removed. You can understand how much time you have before the demand runs out.

### What-If Analysis - Workload Planning: Hyperconverged and VMC on AWS

You can perform Hyperconverged Infrastructure workload planning by adding or removing VMs to VMware vSAN activated clusters and running What-If scenarios. VMware Aria OperationsVMware Cloud Foundation Operations shows you if the proposed workload fits or does not fit in the suggested location. If it fits, the results list the prime target cluster and any additional possible locations. The system also projects time remaining before the workload runs out of resources. .

#### Where You Find What-If Analysis - Workload Planning: Hyperconverged

Click **What-If Analysis** in the left menu under **Capacity**. In the Workload Planning: Hyperconverged tile, click **Add VMS** or **Remove VMS**.

#### How What-If Analysis - Workload Planning: Hyperconverged Works

You define scenarios that can potentially add or remove workloads to VMware vSANenvironment. The workload scenarios are based on VMs associated with specific storage policy related factors (such as FTT, RAID).

**NOTE**
When a workload is added based on imported VMs, and the VM is currently in a VMware vSAN-activated cluster, the VMware vSAN policy settings are not applied and the current VM disk space is taken as is.

## Capacity and Cost Planning Support for Virtual Machines - VMC Datacenter

You can now perform capacity planning and cost calculations for a virtual machine (VM) in hyper-converged environment where the VM is part of the Vmware Cloud on Amazon Web Services (VMC) cluster. VMware Aria OperationsVMware Cloud Foundation Operations provides accurate capacity recommendations and cost calculations when you add or remove VMs in hyper-converged environment from VMC data centers.

The cost calculation is based either on bills collected by VMC adapter or based on reference. To know more about VMC costing, see the topic VMware Cloud on AWS Cost Management in VMware Aria Operations in the *VMware Aria OperationsVMware Cloud Foundation Operations Configuration* guide.

### Add or Remove VMs: Hyperconverged

As part of the What-If workload planning for hyperconverged infrastructure, Workload Planning: Hyperconverged is the pane you use to fill in the details of your virtual machines. You select where to add or remove the workload, configure it yourself or use an existing VM as a template, and establish a time frame. The advanced configuration option lets you define your configuration more precisely.

## Where You Find Workload Planning

From the menu, select **Home** and **Optimize Capacity › What-If Analysis** in the left pane. Click **Add VMS** or **Remove VMS** in the **Workload Planning: Hyperconverged** pane.

**Table 133: Workload Planning: Hyperconverged Add Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario |
| **Location** | Where do you want to add the virtual machines? Select from the list of existing data centers. You can optionally select the exact cluster where you want the virtual machine to reside. |
| **ApplicationProfile/Configure** | Allows you to configure the virtual compute resource, including vCPU, Memory, and Disk Space. |
| **Application Profile/Import Import from existing VM** | Displays the Select VMs dialog box where you can select one or more existing VMs to use as templates for your workload. Once you have made your selections, you return to this screen to enter the quantity of each selected VM you want to incorporate as templates into your workload. |
| Select your workload:<br>• **CPU**<br>• **Memory**<br>• **Disk space** | With the **Configure** radio button selected, you can size your workload by defining values for vCPU, Memory, and Disk Space. These are your allocation values. |
| **Expected Utilization** | Set the projected percentage of total workload capacity you expect to average. Click **Advanced Configuration** to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning. VMware Aria OperationsVMware Cloud Foundation Operations calculates demand based on these values. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| **Annual Projected Growth** | Set the percentage by which you expect your capacity to grow, annually. Click **Advanced Configuration** to set the percentage growth of CPU, Memory, and Disk individually.<br><br>For example, if the utilization is 100 at the start date, and you set the annual growth % to 10%, then at the end of the year the utilization will grow to 110.<br><br>The Annual Projected Growth can be set to 0% if no growth is expected. |
| Number of VMs (optional)/**Quantity** | You can optionally select how many VMs to spread the workload across. |
| Additional vSAN configuration | Configure additional VMware vSAN details such as swap space, host failures to tolerate, fault tolerance method, and Dedup. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. |
| **Run Scenario** | Click to run the scenario. The system calculates whether it fits into the location you selected. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

**Table 134: Workload Planning: Hyperconverged Remove Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario. |
| **Location** | From where do you want to remove the VMs? Select from the list of existing data centers. You can optionally select the exact cluster from where you want to remove the workload. |
| **ApplicationProfile**/**Configure** | Allows you to configure the virtual compute resource, including vCPU, Memory, and Disk Space.<br>After you have configured the scenario, enter the quantity of custom VMs that you want to remove. |
| **Application Profile**/**Import Existing VMs** | Displays the **Select VMs** dialog box where you can select one or more existing VMs. Once you have made your selections, you return to this screen to enter the quantity of each selected VM you want to remove from your workload.<br><br>**NOTE**<br>The recommended limit is 100 VMs as a maximum for workload removal. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br><br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br><br>Click **OK** when done. |
| **Application Profile** / **Custom**: Choose your workload<br>• **CPU**<br>• **Memory** | With the **Configure** radio button selected, you can size your workload by defining values for vCPU, Memory, and Disk Space. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| • **Disk space** | |
| **Expected Utilization** | Set the projected percentage of total workload capacity you expect to average. Click **Advanced Configuration** to set the percentage of expected utilization for CPU, Memory, and Disk individually and to select thin or thick provisioning. |
| Number of VMs (optional)/ **Quantity** | You can optionally select how many VMs to spread the workload across. |
| Additional vSAN configuration | Configure additional VMware vSAN details such as swap space, host failures to tolerate, fault tolerance method, and Dedup. |
| **Start Date/End Date** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. You can also leave the end date blank. |
| **Run Scenario** | Click to run the scenario. The system calculates the impact on the cluster (time remaining and capacity remaining) when removing the workload. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

**Results: Add or Delete VMs to Hyperconverged Infrastructure**

The scenario results are displayed when you run the scenario. In Private Cloud Data Center you can view the recommendation which provides details about the number of VMs to be added or removed from the VMware Cloud. You can also view whether the workload fits in your cloud environment and the cost increase or cost saving depending on whether you are adding or removing a VM from the VMware cloud. The Public Cloud tile displays the cost increase or savings across public clouds like Google Cloud, VMware Cloud on AWS, Amazon Web Services, IBM Cloud and others.

### Scenario Summary Page for Workload Planning: Hyperconverged

The scenario summary page for Workload Planning - Hyperconverged displays the workload values after you add or remove VMs to VMware vSAN activated clusters.

**Summary Section**

The summary section displays the inputs that you provided for the Add VMS or Remove VMS scenario. You can change the workload location by selecting a different VMware vSAN activated clusters. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

**Private Cloud and Public Cloud sections for Add VMS**

The Private Cloud: Datacenter, Public Cloud, and Private Cloud sections display information which help you understand where your workload will fit, the associated costs, and the time remaining based on peak CPU, Memory and Disk Space for demand and allocation model after you add VMs. You can choose to commit the Private Cloud: Datacenter results by clicking the **COMMIT SCENARIO** button. To see more details, click **VIEW DETAILS**.

You can see the cost per month for different private and public clouds. Click **Learn More** for a more detailed comparison based on your workload.

**Scenario Results Section for Remove VMS**

This section displays a graphical view of the projected CPU, Memory and Disk Space values for the demand and allocation models after the VMs are removed. You can understand how much time you have before the demand runs out.

## What-If Analysis - Infrastructure Planning: Traditional

You define scenarios that can potentially add capacity to actual data centers or remove capacity from actual data centers. VMware Aria OperationsVMware Cloud Foundation Operations models the scenario and calculates whether your desired workload can fit in the targeted data center or custom data center.

### Where You Find Infrastructure Planning: Traditional

In the left menu, click **Capacity › What-if Analysis**. The What-if Analysis Plan page opens. Click **Add Hosts** or **Remove Hosts** in the Infrastructure Planning: Traditional tile.

### How the What-If Analysis for Infrastructure Planning: Traditional Works

Infrastructure Planning for traditional environments lets you to forecast successfully the impact of adding capacity to your environment or removing capacity from your environment. By trying various scenarios, you can arrive at an optimum configuration. Once you select the Infrastructure Planning: Traditional pane, you can choose where you want to locate the additional capacity or from where you can remove the existing capacity.

In selecting the profile while removing capacity, you can select a profile only from server types that exist in your cluster.

In selecting the profile while adding capacity, you have two options:

- Select a server type from a list of commercially available servers. You can select from a list of 1) server types already in your cluster or 2) all server types approved for purchase.
- Configure a custom server manually by specifying CPU attributes, memory, and cost.

When you have set the profile for the new server, enter the number of servers to purchase or remove and the start and end date for the period when you want the scenario to be active. The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster. The system can project scenarios ending up to one year from the current date. By default, the starting date is today and the ending date is one year from today.

At this point, you can save the scenario to edit or run later on. A list of saved scenarios is available on the What-If Analysis main page. Otherwise, run the scenario to get the VMware Aria OperationsVMware Cloud Foundation Operations analysis and assessment of your plan.

The system displays immediately the impact on cluster size of the additional or lesser amount of CPU and memory, and shows the total cost of adding or removing the specified capacity. The system also shows whether adding new capacity or removing capacity extends or shrinks the time remaining before CPU or memory runs out.

As well, the system displays a graphic depiction of resource use. For each attribute value - CPU and memory - the amount by which the workload increases or decreases the percentage of total capacity used is shown against a time line.

## Add or Remove Hosts

As part of the What-If analysis for physical infrastructure planning for traditional environments, Infrastructure Planning: Traditional pane is what you use to fill in the details of your What-If scenario. You select where to add or remove hosts, use an existing server type, or configure it yourself (when you add capacity), and establish a time frame.

### Where You Find Physical Infrastructure

At the What-If Analysis screen, click **Add Hosts** or **Remove Hosts** in the Infrastructure Planning: Traditional pane.

**Table 135: Add Hosts Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario |
| **Location** | Where do you want to add capacity? Select from the list of existing data centers, then select the cluster where you want one or more servers to reside. |
| **Server Details** | Clicking **Select Server** displays the Select Server Type dialog box, where you can select a commercial brand server or configure a custom server.<br>Number of Servers to add: increment the Quantity counter up to the number of servers you want. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the What-If scenario. |
| **Run Scenario** | Click to run the scenario. The system calculates the cost of the scenario and determines any new time remaining number. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

The system displays immediately the impact on cluster size of the additional CPU and memory, and shows the total cost of adding the specified capacity. The system also shows in graphical form whether adding the new capacity extends the time remaining before CPU or memory runs out.

**Table 136: Remove Hosts Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario |
| **Location** | From where do you want to remove capacity? Select from the list of existing data centers, then select the cluster from where you want to remove one or more servers. |
| **Server Details** | Clicking **Select Server** displays the Select Server Type dialog box, where you can select only the server types that exist in your selected cluster.<br>The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster. |
| **Start Date/End Date** | Select from pop-up calendars the start and end date for the What-If scenario. You can select to keep the end date blank. |
| **Run Scenario** | Click to run the scenario. The system determines any new time remaining number. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

The system displays the time remaining and the impact on CPU and memory with reduced capacity. The system also shows in graphical form whether removing capacity decreases the time remaining before CPU or memory runs out.

You can also see that the cost is based on the original purchase cost.

### Scenario Summary Page for Infrastructure Planning: Traditional

The scenario summary page for Infrastructure Planning - Traditional displays the capacity values after you add or remove hosts to a server in a data center and cluster you specify.

### Summary Section

The summary section displays the inputs that you provided for the Add Hosts or Remove Hosts scenario. You can change the capacity location by selecting a different vCenter and cluster combination, or server. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

### Scenario Results Section

This section displays a graphical view of the projected CPU, Memory and Disk Space capacity values for the demand and allocation models after the hosts are added or removed. You can understand how much time you have before capacity runs out and the costs savings. Click the **COMMIT SCENARIO** button to reserve this capacity.

## What-If-Analysis - Infrastructure Planning: Hyperconverged

You can perform infrastructure planning by adding or removing Hyperconverged Infrastructure (HCI) nodes in vSAN activated clusters and running What-If scenarios. VMware Aria OperationsVMware Cloud Foundation Operations displays the cost, time remaining, and capacity remaining for CPU, memory, and disk space in the scenario results.

### Where You Find What-If Analysis - Hyperconverged Infrastructure

In the left menu, click **Capacity › What-if Analysis**. The What-if Analysis Plan page opens. Select the Infrastructure Planning: Hyperconverged tile. To run a What-If scenario click **Add HCI Nodes** or **Remove HCI Nodes**.

### How What-If Analysis - Hyperconverged Infrastructure Works

You can add hyperconverged infrastructure to your VMware vSAN activated environment evaluate the increase in HCI capacity and cost. You can add up to 64 hosts per vSAN cluster. This number accounts for existing hosts in the cluster. VMware Aria OperationsVMware Cloud Foundation Operations only lists vSAN and vXRail clusters in the location property. You can select existing server types from these locations and change the number of instances of these servers to add to your scenario.

## Add or Remove HCI Nodes

As part of the what-if analysis for physical infrastructure planning for hyperconverged environments, the Infrastructure Planning: Hyperconverged pane is what you use to fill in the details of your what-if scenario. When you add an HCI node, you can select an existing server type from your vSAN activated data center and change the number of instances of this server to calculate storage, compute capacity, time remaining, and cost. You can run the Remove HCI Nodes scenario to see the capacity changes after you remove HCI nodes from your data center.

### Where You Find Workload Planning

At the **What-If Analysis** page, click **Add HCI Nodes** or **Remove HCI Nodes** in the **Infrastructure Planning: Hyperconverged** pane.

**Table 137: Add HCI Nodes Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario. |
| **Location** | Where do you want to add the HCI node? Select from the list of existing data centers. You must also choose the exact cluster where you want the HCI node to reside. |
| **Server Details** | Allows you to select an existing server type to calculate capacity, time, and storage remaining based on the number of instances of the server. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| **Number of servers to add** | How many instances of the server do you want to add?<br>**NOTE**<br>Only 60 new hosts can be added to the specified vSAN cluster as the maximum allowed is 64 hosts. |
| **Start Date/End Date** | Select from pop-up calendars the start and end date for the workload. The end date cannot be later than one year from the current date. |
| **Run Scenario** | Click to run the scenario. The system calculates whether it fits into the location you selected. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

**Table 138: Remove HCI Nodes Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario. |
| **Location** | From where do you want to remove capacity? Select from the list of existing data centers, then select the cluster from where you want to remove the server(s). |
| **Server Details** | Clicking **Select Server** displays the Select Server Type dialog box, where you can choose only the server types that exist in your selected cluster.<br>The number of servers that you plan to remove is limited by the number of selected server types available in the selected cluster. |
| **Implementation date/End Date (optional)** | Select from pop-up calendars the start and end date for the what-if scenario. You can choose to keep the end date blank. |
| **Run Scenario** | Click to run the scenario. The system determines any new time remaining number. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

### Scenario Summary Page for Infrastructure Planning: Hyperconverged

The scenario summary page for Infrastructure Planning - Hyperconverged displays the capacity values after you add or remove Hyperconverged Infrastructure (HCI) nodes in vSAN activated clusters that you specify.

#### Summary Section

The summary section displays the inputs that you provided for the Add HCI Nodes or Remove HCI Nodes scenario. You can change the vSAN activated cluster location by selecting a different vCenter and cluster combination, or server. Alternatively, you can edit or save the scenario. You can run saved scenarios later, and commutatively.

#### Scenario Results Section

This section displays a graphical view of the projected CPU, Memory and Disk Space capacity values for the demand and allocation models after the HCI nodes are added or removed. You can understand how much time you have before capacity runs out and the costs savings. Click the **COMMIT SCENARIO** button to reserve this capacity.

## What-If-Analysis - Migration Planning: VMware Cloud

What-If-Analysis - Migration Planning, allows you to evaluate your plan for migrating or moving workloads across different VMware clouds. You can compare capacity and cost of workload across Oracle Cloud VMware Solution (OCVS), VMware Cloud on Dell EMC (VMCD), VMware Cloud on Amazon Web Services (AWS), Azure VMware Solution (AVS), and Google Cloud VMware Engine (GCVE). VMware Aria OperationsVMware Cloud Foundation Operations evalautes the migration plan, calculates the cost and capacity requirements, and provides cost estimate for the selected VMC workload.

### Where You Find What-If Analysis - Migration Planning

In the left menu, click **Capacity › What-if Analysis**. The What-if Analysis Plan page opens. In the Migration Planning: VMware Cloud tile, click **Plan Migration**.

### How What-If Analysis - Migration Planning Works

The What-If-Analysis feature of Capacity Optimization lets you forecast successfully the impact of migrating a workload to the VMware Cloud instance such as Oracle Cloud VMware Solution (OCVS), VMware Cloud on Dell EMC (VMCD), VMware Cloud on Amazon Web Services (AWS), Azure VMware Solution (AVS), and Google Cloud VMware Engine (GCVE). Once you select the Migration Planning screen, choose whether you want to run the scenario against a VMware Cloud on AWS or other type of cloud accounts. For the VMware Cloud, select the region where you want to migrate the workload.

When you have set the profile for the migrating workload, run the scenario to get the analysis and assessment of your plan. You can select one VMware cloud at a time and get the estimate for migration planning cost. Alternatively, you can save the scenario to edit or run later on. A list of saved scenarios is available in the Saved Scenarios tab on the What-If analysis page.

If you selected VMware Cloud on AWS for your scenario, the results list the VMware Cloud on AWS Assessment, with details of the VMware configuration. The result also displays the resource-use-level cost and the monthly purchase cost for an on-demand subscription. In addition, the result displays the resource-use-level cost and monthly purchase cost for one-year and three-year subscriptions.

### About Clouds

The system might provide a recommendation based on the cost of placing the workload on different VMware clouds. This cost-based recommendation varies for different clouds.

For VMware Cloud on AWS, the system displays the resource-use-level cost and the monthly purchase cost for an on-demand subscription, plus those same costs for one-year and three-year subscriptions.

VMware cloud costs are based on the selected configuration, that is, the allocated resources.

## Migration Planning: VMware Cloud

As part of the What-If Analysis function, Migrate is the form you use to fill in the details of your what-if scenario. You choose where to migrate the workload, then select the region.

### Where You Find Migration Planning

At the What-If Analysis screen, click **Plan Migration** in the **Migration Planning: VMware Cloud** tile.

When you run a scenario for What If: Migration for VMware Cloud, VMware Aria OperationsVMware Cloud Foundation Operations might suggest the right cloud Instance suitable for the Workload Configuration selected by you. VMware Aria OperationsVMware Cloud Foundation Operations also calculates the cost for that VMware Cloud's instance and displays the same.

**Table 139: Migrate Options**

| Option | Description |
|---|---|
| **SCENARIO NAME** | Name of your scenario |
| **SELECT CLOUDS** | Where do you want to migrate the workload?<br>Options:<br>• **VMware Cloud on AWS**<br>• **Azure VMware Solution (AVS)**<br>• **Google Cloud VMware Engine (GCVE)**<br>• **Oracle Cloud VMware Solution (OCVS)**<br>• **VMware Cloud on Dell EMC (VMCD)**<br>**NOTE**<br>You can select regions for the above cloud options. |
| **CLUSTER SETTINGS** | Specify the following cluster details:<br>• Enter the **Instance Type**.<br>• Enter the **Slack Space** in percentage.<br>• Enter the **Steady state CPU headroom** in percentage. |
| **APPLICATIONPROFILE/Configure** | Using the Application Profile you can configure the virtual compute resources, like vCPU, memory, and storage. |
| Select Your Workload:<br>• **CPU**<br>• **Memory**<br>• **Disk Space** | With the Configure radio button selected, you can size your migrating workload by defining values for vCPU, memory, and storage. |
| **Expected Utilization** | Specify the expected utilization or Click **ADVANCED CONFIGURATION** and specify the values for the following:<br>• **CPU**<br>• **Memory**<br>• **Disk Space**<br>• **Disk Space Provisioning** - Select either Thin or Thick. |
| **Annual Projected Growth** | Specify the annual growth rate so that the system adjust the scenario calculations or click **ADVANCED CONFIGURATION** and specify the values for the following:<br>• **CPU**<br>• **Memory**<br>• **Disk Space** |
| **Number of VMs (OPTIONAL)** | You can optionally choose how many VMs to spread the workload across. |
| **Additional vSAN configuration** | Select **Account for Swap Space** to reserve swap space for any unreserved virtual machine memory.<br>• Select the **Host failures to tolerate value** from the drop-down list.<br>• Select the **Fault Tolerance Method**, the options are RAID -1 and RAID-5.<br>• Select the **Dedup** value from the drop down list. |
| **APPLICATION PROFILE/Import from existing VM** | Displays the **Select VMs** button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br><br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br><br>Click **OK** when done. |
| **RUN SCENARIO** | Click to run the scenario. The system calculates whether it fits into the location you chose. |
| **SAVE** | **SAVE** the scenario. |
| **CANCEL** | **CANCEL** the scenario. |

**VMware Cloud on AWS Assessment - Results**

The scenario results are displayed when you run the scenario. You can view the recommendation which provides details about the number of hosts required for the VMware Cloud. You can also view the Total Cost associated with the recommended VMware Cloud for 3 years subscription and the Total Capacity Usage details for CPU, memory and disk space.

For VMware Cloud on AWS Assessment, you can edit the following options.

- **Edit Configuration** - you can edit the change in Reserved Capacity CPU, Reserved Capacity Memory, Fault Tolerance, and RAID Level values and save the values to the original configuration.
- **Change Plan** - you can use the **Choose Plan** option to change your subscription plan, the available options are one-year plan, three-year plan, or Pay-As-You-Go.
- **Edit Discount** - you can use the edit discount option to specify the discount value, the total cost for the subscription is equal to the actual utilization cost minus the discount percentage.

**What-If-Analysis - Migration Planning: Public Cloud**

You define scenarios that can potentially migrate workloads to a public cloud. Use this scenario to determine where to move the workloads. VMware Aria OperationsVMware Cloud Foundation Operations models the scenario and calculates the cost and capacity to fit your desired workload.

**Where You Find Migration Planning: Public Cloud**

In the left menu, click **Capacity › What-if Analysis**. The What-if Analysis Plan page opens. In the **New** tab, click the **PLAN MIGRATION** button in the **Migration Planning: Public Cloud** card.

**How What-If Analysis - Migration Planning Works**

This feature of Capacity Optimization lets you to forecast successfully the impact of migrating a workload to a public cloud instance such as Amazon Web Services, IBM Cloud, Microsoft Azure, or Google Cloud. Slect the region where you want to migrate the workload. If the public clouds listed out of the box do not suit your needs, you can also define your own public cloud and upload a rate card.

In defining the profile of your workload, you have two options:

- Configure the workload manually by specifying vCPUs, memory, storage, and expected use percentage.

- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system allows you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for the migrating workload, run the scenario to get the VMware Aria OperationsVMware Cloud Foundation Operations analysis and assessment of your plan. You can also select up to three public clouds to compare results. Alternatively, you can save the scenario to edit or run later on. A list of saved scenarios is available in the **Saved Scenarios** tab on the What-If analysis page.

For a public cloud target, the system lets you know immediately if the workload proposed for migration fits or does not fit in the suggested location. For example, if you selected AWS and the workload fits, the results list the Amazon Web Services Assessment, with details of the VMware Configuration and the AWS Equivalent. If the proposed workload does not fit, an error message appears: "Unable to identify a matching configuration instance in target location."

**About Clouds**

The system might provide a recommendation based on the cost of placing the workload on different clouds. This cost-based recommendation varies for different clouds. You can modify the costs for public clouds by uploading a new rate card.

Public cloud costs are based on the selected configuration, that is, the allocated resources.

The public instance is selected based on the close proximity rule, with simulated resource allocation values. In some scenarios, an exact configuration match is not available in the list. Due to this lack of availability, the public cost can be inherently higher in comparison.

## Migration Planning: Public Cloud

As part of the What-If Analysis function, Migrate is the form you use to fill in the details of your what-if scenario. You choose where to migrate the workload, then select the region.

**Where You Find Migration Planning**

At the What-If Analysis screen, click **PLAN MIGRATION** in the Migration Planning: Public Cloud card.

When you run a scenario for What If: Migration for Public Clouds (Not VMC), VMware Aria OperationsVMware Cloud Foundation Operations might suggest the Public Cloud Instance suitable for the Workload Configuration selected by you. VMware Aria OperationsVMware Cloud Foundation Operations also calculates the cost for that Public Cloud's instance and displays the same.

**Table 140: Migrate Options**

| Option | Description |
|---|---|
| **SCENARIO NAME** | Name of your scenario |
| **SELECT CLOUDS** | Where do you want to migrate the workload?<br>Options:<br>• Aamazon Web Services<br>• IBM Cloud<br>• Microsoft Azure<br>• Google Cloud |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>The cloud providers added in the Add Cloud Provider page are also included in the list.<br><br>You can select a maximum of three public clouds at a time for comparison. Hold the **Shift** key to select more than one public cloud provider. |
| **ADD CLOUD PROVIDERS** | You can add or edit the cloud providers and also edit the rate card of each individual cloud provider. |
| **APPLICATIONPROFILE/Configure** | Using the Application Profile you can configure the virtual compute resources, like vCPU, memory, and storage. |
| Select Your Workload:<br>• **CPU**<br>• **Memory**<br>• **Disc Space** | With the Configure radio button selected, you can size your migrating workload by defining values for vCPU, memory, and storage. |
| **APPLICATION PROFILE/Import from existing VM** | Displays the Select VMs button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group.<br>Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br><br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br><br>Click **OK** when done. |
| Number of VMs (OPTIONAL)/**Quantity** | You can optionally choose how many VMs to spread the workload across. |
| **RUN SCENARIO** | Click to run the scenario. The system calculates whether it fits into the location you chose. |
| **SAVE** | **SAVE** the scenario. |
| **CANCEL** | **CANCEL** the scenario. |

**What-If Analysis - Data Center Comparison**

You can select virtual machines to determine which of the preferred data centers (along with a specific choice of cluster or default cheapest cluster) are best fit from both cost effectiveness and capacity requirements perspective. The comparison helps you to find the right data center to place the workload from cost and capacity perspective.

**Where You Find What-If Analysis – Data Center Comparison**

In the left menu, click **Capacity › What-if Analysis**. The What-if Analysis Plan page opens. In the **New** tab, click **COMPARE DATACENTERS** in the Datacenter Comparison: Private Cloud tile.

## How What-If Analysis - data center Comparison Works

This feature of Capacity Optimization lets you to compare cost across data centers within the private cloud environment. After you select the Datacenter Comparison screen, choose one or more data centers to compare the cost and run the scenario. VMware Aria OperationsVMware Cloud Foundation Operations suggests which data center is most cost effective for the selected workload.

In defining the profile of your workload, you have two options:

- Configure the workload manually by specifying CPU, memory, disk space, expected utilization, and annual projected growth.
- Use an existing VM or VMs as templates, importing all the attributes of the selected VMs to your workload scenario. The system allows you to specify how many copies of each selected VM you want to add to the proposed workload.

When you have set the profile for comparing the workload, run the scenario to get the VMware Aria OperationsVMware Cloud Foundation Operations analysis and assessment of your plan. You can select up to three data centers to compare results. Alternatively, you can save the scenario to edit or run later. A list of saved scenarios is available in the Saved Scenarios tab on the What-If analysis page.

Cost varies from one data center to another depending on cost settings, which include cost drivers such as servers, facility, power, labor, license, network, and storage.

The data center comparison feature solves this problem by allowing you to select a data center which suits your requirement, is least expensive, and has adequate capacity.

### Datacenter Comparison

As part of the What-If Analysis function, Compare Datacenters is the form you use to fill in the details of your What-If scenario. Use this scenario to compare cost across data centers within the private cloud environment.

## Where You Find Compare Datacenters

At the **What-If Analysis** page, click **Compare Datacenters** in the pane titled Datacenter Comparison.

**Table 141: Compare Datacenter Options**

| Option | Description |
|---|---|
| **Scenario Name** | Name of your scenario. |
| **Select Datacenters** | Select the data centers for which you want to compare the costs. Use the advanced filter search to find a datacenter. |
| **Application Profile/Configure** | Using the Application Profile, you can configure the virtual compute resources, like CPU, memory, disk space, expected utilization, and annual projected growth. |
| Select Your Workload:<br>• **CPU**<br>• **Memory**<br>• **Disk Space**<br>• **Expected Utilization**<br>• **Annual Projected Growth** | With the Configure radio button selected, you can size your workload by defining values for CPU, memory, disk space, expected utilization, and annual projected growth. |
| **Application Profile/Import from existing VM** | Displays the Select VMs button. When selected, displays the Select VMs workspace, where you can choose one or more existing VMs to use as templates for your workload. You can filter VMs by name, tags, vCenter Server, or custom group. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Once you have made your selections, you return to this screen to enter the quantity of each chosen VM you want to incorporate as templates into your workload. |
| **Import from Custom Profile** | Click the **SELECT CUSTOM PROFILES** button to select one or more Custom Profiles to perform comparative analysis of capacity and cost across different datacenters and clusters in your private cloud.<br><br>In the Select Custom Profiles dialog box which opens, select from a list of custom profiles and click the the **>** icon to add the custom profile to the **SELECTED** list.<br><br>Click **OK** when done. |
| Number of VMs (OPTIONAL)/**Quantity** | You can optionally choose how many VMs to spread the workload across. |
| **Implementation date/End Date (optional)** | You can specify the Start Date and End Date to compute the data center infrastructure cost for a specific time period. |
| **Run Scenario** | Click to run the scenario. The system calculates the cost of migration and checks whether the selected workload fits into the location you have chosen. |
| **Save** | **Save** the scenario. |
| **Cancel** | **Cancel** the scenario. |

**Retain Historical Data of VMs Migrated Using VMware Hybrid Cloud Extension**

In VMware Aria OperationsVMware Cloud Foundation Operations you can use VMware Hybrid Cloud Extension (HCX) to perform application migration, workload rebalancing, and business continuity across data centers and clouds. You can also migrate workloads from on-premises data centers to VMware Cloud.

Earlier when you performed HCX bulk migration to migrate workloads from one datacenter to another datacenter, or from one datacenter to VMware Cloud. VMware Aria OperationsVMware Cloud Foundation Operations failed to retain the historic metrics.

Now VMware Aria OperationsVMware Cloud Foundation Operations has implemented a solution which is triggered during HCX migration. The event helps VMware Aria OperationsVMware Cloud Foundation Operations to collect details and manage the target vCenter to identify the workload migration.

After mapping the right attributes of VMs in the source datacenter with the VMs in the destination datacenter, you can verify if VMware Aria OperationsVMware Cloud Foundation Operations is able to retain all the historical metrics. The HCX migration types supported in VMware Aria OperationsVMware Cloud Foundation Operations are:

- Bulk Migration
- vMotion Based Migration (hot & cold)
- Replication Assisted Migration

HCX vMotion

VMware Aria OperationsVMware Cloud Foundation Operations performs the following actions during HCX vMotion.

- Get the target `VCI`, `VM-VC-MOID` from event using the resource key.
- Get the source `VCID`, `VM-VC-MOID` from event using the resource key.
- Maps the correct target `VCID`, `VM-VC-MOID` to the source VM in VMware Aria OperationsVMware Cloud Foundation Operations.

vMotion Generic Scenario

VMware Aria OperationsVMware Cloud Foundation Operations performs the following actions during vMotion.

- Get the target `VCI`, `VM-VC-MOID` from event using the resource key.
- Detects the discovered target VM in VMware Aria OperationsVMware Cloud Foundation Operations , based on the `VCID`, `VM-VC-MOID` attribute.
- For the detected target VMs, get VMs with **VM Entity Instance UUID**, and map the `VCID`, `VM-VC-MOID` for these VMs.
- Search `VCID`, `VM-VC-MOID` in the event message to find the actual vMotion VM.
- Set the correct target `VCID`, `VM-VC-MOID` to the right VM in on source VMware Aria OperationsVMware Cloud Foundation Operations.

**NOTE**
To know more about HCX Migration, see VMware HCX Product Documentation.

## Configuring Cost

SDDC costing is out-of-the box with VMware Aria OperationsVMware Cloud Foundation Operations. There is no integration required with vRealize Business for Cloud.

## Cost Overview

VMware Aria OperationsVMware Cloud Foundation Operations now supports costing for private clouds, public clouds, and VMware Cloud Infrastructure. You can track expenses for a single virtual machine (VM), and how these expenses attribute to the overall cost associated with your private cloud accounts and VMware Cloud Infrastructure accounts.

The Cost Overview home page provides all the details about the costs associated with your VMware Cloud Infrastructure accounts, public cloud accounts, and your private cloud accounts. You can view the Total Cost of Ownership, Potential Savings, and Realized Savings for your VMware Cloud Infrastructure cloud accounts and vSphere Private Cloud accounts, and Total Cost of Ownership for your private cloud accounts.

You can view the cost details for the following private and public cloud accounts in VMware Aria OperationsVMware Cloud Foundation Operations.

- vSphere On-Prem
- VMware Cloud Foundation (VCF)
- VMware Cloud on AWS
- Azure VMware Solutions (AVS)
- Google Cloud

- All the values shown in vSphere On-Prem page might not match with the respective metrics at vSphere world level. That is because the metric at vSphere world level has values aggregated for all clouds like, Private, VMC, AVS, VMC-D, GCVE, and OCVS. So the values shown in Overview page are obtained after subtracting the aggregated metric values of VCF clouds from the metric value at vSphere World object.

**Private Cloud - Example: vSphere On-Prem**

The cost component of vSphere On-Prem private cloud account and VCF cloud account are Total Cost of Ownership, Potential Savings, and Realized Savings.

**Total Cost of Ownership** - The total cost of ownership widget displays the cost expenditure by capacity, by cost drivers and by datacenter. You can use the by capacity pie chart to view the compute, storage, and VM direct cost associated with your VMware Cloud Infrastructure cloud accounts. The cost drivers bar graph provides details of the cost drivers associated with your VCF cloud accounts and the horizontal graph for datacenters provides the expense details of your VMware Cloud Infrastructure cloud accounts for individual datacenters.

**NOTE**
The cost displayed in the Total Cost of Ownership widget might not match with the TCO metric at vSphere world level since it is the Total Aggregated Cost.

All the values shown in vSphere On-Prem widget might not match with the respective metrics at vSphere world level. This is because the metric at vSphere world level has values aggregated for all the clouds, like Private, VMC, AVS,VMC-D, GCVE, and OCVS. Hence the values shown in Overview page are obtained after subtracting the aggregated metric values of VMware Cloud Infrastructure from the metric value at vSphere World object.

**Potential Savings**

The potential savings widget displays the amount of savings you can potentially make from your VMware Cloud Infrastructure cloud accounts and vSphere Private cloud accounts. The pie chart for resources displays the cost savings opportunity across distributed across Idle VMs, Orphaned VMs, Oversized Hosts, Powered Off VMs, Reclaimable Hosts, and VM Snapshots. The horizontal graph for datacenters provides the overall potential cost savings for your cloud infrastructure and potential savings for individual datacenters. To know more about Potential Savings see, Potential Cost Savings Dashboard

**NOTE**
The potential savings option is not available for public cloud accounts.

**Realized Savings**

The realized savings widget displays the amount of savings you can potentially make from your VMware Cloud Infrastructure cloud accounts and vSphere Private cloud accounts. The pie chart for resources displays the cost savings opportunity distributed across Idle VMs, Orphaned VMs, Oversized Hosts, Powered Off VMs, Reclaimable Hosts, and VM Snapshots. The horizontal graph for datacenters provides the realized cost saving for your overall cloud infrastructure and realized savings for individual data centers. To know more about Potential Savings see, Realized Cost Savings Dashboard. To know more about reclamation cost savings, see Realized Cost Savings Using Reclamation Suggestion.

**NOTE**
The realized savings option is not available for public cloud accounts.

**Dashboards**

The cost dashboards widget lets you compare the cost of VMware Cloud Infrastructure with other public cloud platforms. You can analyze the cloud comparison results and identify the opportunities to manage your cloud resources efficiently. You can click dashboard links and navigate to the respective dashboard from the cost overview page.

**NOTE**
The data displayed in the dashboards might not be specific to the selected cloud type. The dashboards might contain data from all vSphere instances and VMware cloud instances.

**Public Cloud - Example: Google Cloud Platform**

The cost component of public cloud account includes the cost of ownership associated with your public cloud account. The cost components of public cloud account are distributed across accounts, regions, and services. You can select individual account and view the cost associated with that account, region wise or service wise. For Google Cloud Platform the services cost component is replaced by product category.

**NOTE**
The data for the public cloud accounts like Google Cloud is collected using the Cloud Health adapter, if the data is not displayed for this section, you have to deploy and configure the cloud adapter.

To know the granular cost visibility and to track your expenses of virtual machines accurately in a private cloud, see Overview of Cost Drivers.

To know the expenses related to the CPU, memory, and storage for a single virtual machine (VM), and how they attribute to the overall cost associated with your cloud infrastructure, see VMware Cloud on AWS Cost Management in .

To know more about Reference based costing for Google Cloud VMware Engine (GCVE) and Azure VMware Solution (AVS), see Reference Based Costing for / .

**VMware Cloud on AWS Cost Management in VMware Aria OperationsVMware Cloud Foundation Operations**

IT teams spend on purchasing infrastructure from Vmware Cloud on AWS (VMC). Now they can transfer these expenses (CPU, Memory, and Storage) to the application teams using VMC cost allocation. The cost allocation mechanism lets you view the expenses related to the CPU, memory and storage for a single virtual machine (VM), this helps you to determine the overall cost associated with your cloud infrastructure.

To use the VMC costing feature you must set the **Billing Enable** option in **Advance Settings** section of a VMC adapter to true. If it is set to false, the costing is based on the reference cost.

**VMC Costing - Points to Remember**

- The bill expenses or reference based costs are divided into CPU : memory : storage ratios, you can edit ratio, region, and discount from the Cost Settings for Financial Accounting Model topic.
- The bill expenses are allocated to clusters based on the region to which the cluster belongs.

  **NOTE**
  Some of the bill expenses (co-related to component resource objects in VMware Aria OperationsVMware Cloud Foundation Operations) are divided across all the clusters, since at present VMware Aria OperationsVMware Cloud Foundation Operations does not have an understanding of all the types of expenses.

- If the VMC bills currency format is different from VMware Aria OperationsVMware Cloud Foundation Operations currency format, then the VMC are converted to VMware Aria OperationsVMware Cloud Foundation Operations currency format and published on clusters and VMs. You can find the conversion factor as a property under VMC Organization resource objects.
- The reference based costs that are picked are always on-demand. For example, add VMC vCenter directly to VMware Aria OperationsVMware Cloud Foundation Operations. For VMC, if you set the cloud type as VMware on AWS, then the reference costs of US east (N. Virginia) is picked by default.

The following are some important points to consider when you select reference based costing and bill based costing.

- In case of reference based costing, we consider the Host as Production host and host type as On Demand, and get the base rates for cost Allocation. Even if the host type is Subscription based, we still do costing treating it as On Demand Host Type.
- When you have some unconfigured SDDCs in the organization, VMware Aria Operations might not list all the hosts in the organization. So, if you use bill based costing which uses the list of hosts to calculate the cost, we might not be able to calculate the correct base rates.
- Expenses from the bills of your VMC are distributed using a fair allocation algorithm to CPU, memory, and storage at the VM level. For accurate cost numbers, all the SDDCs must be configured in the given Organization.
- Ability to carry out workload planning with VMC as the destination cloud using the new calculated base rates, based on your bills.

**How Does VMC Cost Allocation Work**

The VMC cost allocation works as per the following sequence of events defined in VMware Aria OperationsVMware Cloud Foundation Operations.

- Discover inventory of VMC using vCenter and VMC adapters.
- Acquire bills for VMC from VMware Cloud Services Platform (CSP) using the VMC native adapter.
- Identify the expenses per cluster using approximate values.
- Using the Total Cost Value, determine CPU, Memory, and Storage base rates.

- Apply base rates on VMs for allocation or utilization depending on the capacity model.

**Google Cloud VMware Engine Cost Management**

IT teams spend on purchasing infrastructure from Google Cloud VMware Engine. Now they can transfer these expenses (CPU, Memory, and Storage) to the application teams using Google Cloud VMware Engine cost allocation. The cost allocation mechanism lets you view the expenses related to the CPU, memory and storage for a single virtual machine (VM), this helps you to determine the overall cost associated with your cloud infrastructure.

**Google Cloud VMware Engine - Points to Remember**

- To use the Google Cloud VMware Engine costing feature you must set the **Billing Enabled** option in the Advanced Settings section of a Google Cloud VMware Engine adapter to **true**. If it is set to **false**, the costing is based on the reference cost. Similarly, if any project of the Google Cloud VMware Engine adapter instance is configured without the CSP Refresh token, then costing is based on reference cost. For more details about the Billing Enabled and CSP Refresh token options, see Configuring a Google Cloud VMware Engine Instance in VMware Aria Operations..
- The GCVEBill object type is associated with multiple projects of the Google Cloud VMware Engine adapter instance. All projects under the GCVEBill object type need to be configured in VMware Aria Operations to ensure accurate costing.
- If a project does not have any GCVEBill object type associated with the adapter instance, then costing is based on reference costing.
- You can create and use a single service account JSON that is common, similar to a super user account, for all the projects. For more details about configuring single service account JSON, see Configuring a Google Cloud VMware Engine Instance in VMware Aria Operations.
- The bill expenses or reference based costs are divided into CPU : memory : storage ratios, you can edit ratio, region, and discount from the Cost Settings for Financial Accounting Model topic.
- The bill expenses are allocated to clusters based on the region to which the cluster belongs.

  **NOTE**
  Some of the bill expenses (co-related to component resource objects in VMware Aria OperationsVMware Cloud Foundation Operations) are divided across all the clusters, since at present VMware Aria OperationsVMware Cloud Foundation Operations does not have an understanding of all the types of expenses.

- If the Google Cloud VMware Engine bills currency format is different from VMware Aria OperationsVMware Cloud Foundation Operations currency format, then the Google Cloud VMware Engine bills are converted to VMware Aria OperationsVMware Cloud Foundation Operations currency format and published on clusters and VMs. You can find the conversion factor as a property under Google Cloud VMware Engine adapter instance resource objects.
- The reference based costs that are picked are always on-demand. For example, add Google Cloud VMware Engine vCenter directly to VMware Aria OperationsVMware Cloud Foundation Operations. For Google Cloud VMware Engine, if you set the cloud type as Google Cloud VMware Engine, then the reference costs of US east (N. Virginia) is picked by default.

The following are some important points to consider when you select reference based costing and bill based costing.

- In case of reference based costing, we consider the Host as Production host and host type as On Demand, and get the base rates for cost Allocation. Even if the host type is Subscription based, we still do costing treating it as On Demand Host Type.
- When you have some unconfigured Private Clouds in the organization, VMware Aria Operations might not list all the hosts in the organization. So, if you use bill based costing which uses the list of hosts to calculate the cost, we might not be able to calculate the correct base rates.
- Expenses from the bills of your Google Cloud VMware Engine are distributed using a fair allocation algorithm to CPU, memory, and storage at the VM level. For accurate cost numbers, all the Private Clouds must be configured in the given Organization.

- Ability to carry out workload planning with Google Cloud VMware Engine as the destination cloud using the new calculated base rates, based on your bills.

**How Does Google Cloud VMware Engine Cost Allocation Work**

Google Cloud VMware Engine cost allocation works as per the following sequence of events defined in VMware Aria OperationsVMware Cloud Foundation Operations.

- Discover inventory of Google Cloud VMware Engine using vCenter and Google Cloud VMware Engine adapters.
- Acquire bills for Google Cloud VMware Engine from VMware Cloud Services Platform (CSP) using the Google Cloud VMware Engine native adapter.
- Identify the expenses per cluster using approximate values.
- Using the Total Cost Value, determine CPU, Memory, and Storage base rates.
- Apply base rates on VMs for allocation or utilization depending on the capacity model.

**Reference Based Costing for Azure VMware Solution/ Oracle Cloud VMware Solution**

IT teams spend on purchasing infrastructure from Azure VMware Solution (AVS) and Oracle Cloud VMware Solution. Now you can transfer these expenses (CPU, Memory, and Storage) to the application teams using reference-based cost allocation. The cost allocation mechanism lets you view the expenses related to the CPU, memory, and storage for a single virtual machine (VM), this helps you to determine the overall cost associated with your cloud infrastructure.

**Reference Based Costing - Points to Remember**

- The reference-based costs are divided into CPU : memory : storage ratios, you can edit ratio,region, and discount from the Cost Settings for Financial Accounting Model topic.
- The reference-based costs are allocated to clusters based on the region to which the cluster belongs.

The following are some important points to consider when you select reference-based costing.

- When you have some unconfigured private clouds in the organization, VMware Aria OperationsVMware Cloud Foundation Operations might not list all the hosts in the organization.
- Ability to carry out workload planning with Azure VMware Solution/ Oracle Cloud VMware Solution as the destination cloud using the new calculated base rates, based on your bills.
- For reference-based costing, we consider the Host as Production host and host type as On Demand and get the base rates for cost Allocation. Even if the host type is Subscription based, we still do costing treating it as On Demand Host Type.

## Cost Analysis

You can run cost, price, and VMware Cloud Bills analysis for your custom objects and groups. You can compare cost and price metrics across various custom groups. You can run an analysis and then view a list of the top and bottom 5, 10, and 20 objects based on selected metrics. Additionally you can also compare the metrics of cost, price, and VMware Cloud bills across objects.

**Analysis of Cost Metrics**

Cost is what you spend in currency value on managing and maintaining your infrastructure. Total Cost gives you an idea about the overall cost of ownership of your infrastructure. You can analyze cost metrics like Total Cost, Potential Savings, Realized Savings and so on, for your infrastructure that includes, VMs, hosts, and clusters. You can use lists and compare between selected objects types. For example, you can view the Top 10 VMs within a datacenter based on Total Cost for the last month or you can compare Cost Center A, Cost Center B, and Cost Center C based on Total Cost for the past 6 months.

**How Cost Analysis Helps**

Some use cases for cost analysis are as follows:

- Financial reporting helps IT teams justify the expenses they incur to the finance team. It helps track expenses across private clouds, VMware Cloud on AWS, and public clouds.
- To drive accountability for resource usage, you can view the cost of infrastructure at the VM level. You can generate a bill for your usage and drive accountability within the application teams.
- Your infrastructure operations teams may want to quantify savings through reclamation or revoking of resources like idle VMs, snapshots, and so on. In VMware Aria Operations you can see such resources and view the price associated with them. Reclaimable or right sizeable resources have a cost associated that can be seen as Potential Savings.
- Realized Savings covers the cost savings from reclamation opportunities recommended by VMware Aria Operations.

### Metrics and Objects You Can Run Cost Analysis For

Objects for Cost metrics (Total Cost, Realized Saving, Potential Saving, MTD Cost, and Monthly Projected Total Cost):

- vCenter (all objects that have Cost: VM, Host, Cluster, Data Center, Resource Pools, and Datastore)
- CloudHealth (all key objects that have Cost: Services, Accounts, Regions, and so on)
- Horizon (all key objects that have Cost or are introducing Cost in this release – User, Session, VDI Pool, and RDS Farm)
- Container (Business Apps and Custom Groups)
  **Note:** For Custom Groups, it is the summary of Total Cost of all VMs under that custom group.

### Where You Find Cost Analysis

In the left menu, click **Capacity › Cost**, and then click the **Analysis** tab. The **Cost Analysis** page opens on the right. Click **Add** to create an analysis. For more information about creating an analysis, see New or Edit Cost Analysis.

## Analysis of Price Metrics

You can analyze price metrics for your infrastructure. The Total Price metric is based on the pricing policy defined by the user. You can use lists and compare pricing between selected objects types.

### How Price Analysis Helps

Some use cases for price analysis are as follows:

- Typically, when your IT teams want the system to do the metering, but want to charge the application teams using their own Business or Pricing model, they can setup their own rate cards in VMware Aria Operations and in turn use chargeback reports as a billing mechanism for the users to help them understand what they are being charged for.
- Service providers as well as organizations following a chargeback model to drive accountability and to help generate usage/allocation based bills for their application teams or cost centers.
- Chargeback helps IT teams to drive behavior that is better achieved when there is a cost associated with the provisioned resources.

### Metrics and Objects You Can Run Price Analysis For

Objects for Price Metrics (Total Price):

- vCenter (all objects that have Price: VM and Cluster)
- vCloud Director (all key objects that have Price – OrgVdc, Org, vApp, and VM)
- Horizon (all key objects that have Price – User, Session, VDI Pool, and RDS Farm)
- Container (Business Apps and Custom Groups)
  **Note:** For Custom Groups, it is the summary of Total Cost of all VMs under that custom group.

### Where You Find Price Analysis

In the left menu, click **Capacity › Cost**, and then click the **Analysis** tab. Click **Add** to create an analysis. For more information about creating an analysis, see New or Edit Cost Analysis.

### Analysis of VMware Cloud Bills

VMware Cloud bills is the exact value you pay against the VMware Cloud infrastructure through Google Cloud VMware Engine, and VMware Cloud on AWS. Currently these bills are available if the seller on record is VMware. Using this analysis you can view, list, and compare the different bills you have received for the given cloud over a period of time

#### Metrics and Objects You Can Run VMware Cloud Bills Analysis For

Objects for VMware Cloud Bills (Total Component Expense, Outstanding Expense, Monthly Commit Expense, Monthly On demand expense, and Monthly Total Expense):

- VMware Cloud on AWS (all bill objects)
- Google Cloud VMware Engine (all bill objects)

#### Where You Find VMware Bills Analysis

In the left menu, click **Capacity › Cost**, and then click the **Analysis** tab. Click **Add** to create an analysis. For more information about creating an analysis, see New or Edit Cost Analysis.

### Cost Analysis - Metrics and Metric Keys

As part of Cost Analysis you can analyze Total Cost, Realized Saving, Potential Saving, MTD Cost, Monthly Projected Total Cost, Total Price, Total Component Expense, Outstanding Expense, Monthly Commit Expense, Monthly OnDemand Expense, and Monthly Total Expense. Metrics are divided into Cost, Price and Bill sections. The tables below map the metric names as per the cost UI with the metric names and the metric keys for each adapter and analysis (cost, price, and bills).

#### vCenter Adapter Metrics: Analyzing Cost

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Datacenter** | | |
| Total Cost | Cost\|Monthly Aggregated Total Cost (US$/Month) | cost\|aggrTotalCost |
| Realized Saving | Cost\|Realized Savings\|Total Realized Cost (US$) | cost\|realized_savings\|realizedTotalCost |
| Potential Saving (Current month) whole month | Reclaimable\|Potential Savings (US$) | reclaimable\|cost |
| **Object Type: Cluster Compute Resource** | | |
| Total Cost | Cost\|Monthly Cluster Total Cost (US$/Month) | cost\|totalCost |
| **Object Type: Host System** | | |
| Total Cost | Cost\|Monthly Server Fully Loaded Cost (US$/Month) | cost\|totalLoadedCost |
| MTD Cost | Cost\|MTD Server Total Cost (US$) | cost\|totalMTDCost |
| **Object Type: Virtual Machine** | | |
| Total Cost | Cost\|Effective MTD Cost (US$) | cost\|effectiveTotalCost |
| Monthly Projected Total Cost (current month) whole month | Cost\|Monthly Effective Projected Total Cost (US$/Month) | cost\|effectiveProjectedTotalCost |
| Potential Saving | Cost\|Potential Savings (US$) | cost\|reclaimableCost |
| **Object Type: Datastore** | | |
| Total Cost | Monthly Total Cost | cost\|totalCost |

**vCenter Adapter Metrics: Analyzing Price**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Virtual Machine** | | |
| Total Price | Summary\|Metering\|Total price | summary\|metering\|value |
| **Object Type: Cluster Compute Resource** | | |
| Total Price | Summary\|Metering\|Total price | summary\|metering\|value |

**Kubernetes Adapter Metrics: Analyzing Cost**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Kubernetes Node** | | |
| Total Cost | Cost\|Daily Total Cost | cost\|totalCost |
| **Object Type: Kubernetes Namespace** | | |
| Total Cost | Cost\|Daily Total Cost | cost\|totalCost |
| **Object Type: Kubernetes Cluster** | | |
| Total Cost | Cost\|Daily Total Cost | cost\|totalCost |

**VMware Aria Operations for Networks Adapter Metrics: Analyzing Cost**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Operations for Network Application** | | |
| Total Cost | Cost\|Effective MTD Cost | cost\|effective_total_cost |
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | cost\|effective_projected_total_cost |
| **Object Type: Operations for Network Tier** | | |
| Total Cost | Cost\|Effective MTD Cost | cost\|effective_total_cost |
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | cost\|effective_projected_total_cost |

**VMware Aria Automation Adapter Metrics: Analyzing Cost**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Project** | | |
| Total Cost | Cost\|Effective MTD Cost | cost\|effectiveTotalCost |
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | cost\|effectiveProjectedTotalCost |
| **Object Type: Deployment** | | |
| Total Cost | Cost\|Effective MTD Cost | cost\|effectiveTotalCost |
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | cost\|effectiveProjectedTotalCost |

**VMware Aria Automation Adapter Metrics: Analyzing Price**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Project** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |
| **Object Type: Deployment** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |

**VMware Cloud on AWS and Google Cloud VMware Engine metrics : Analyzing Bills**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Component** | | |
| Total Component Expense | Cost\|Total Component Expense | cost\|total_component_expense |
| **Object Type: VmcBill** | | |
| Outstanding Expense | Cost\|Outstanding Expense | cost\|outstanding_expense |
| Monthly Commit Expense | Cost\|Monthly Commit Expense | cost\|monthly_commit_expense |
| Monthly OnDemand Expense | Cost\|Monthly OnDemand Expense | cost\|monthly_on_demand_expense |
| Monthly Total Expense | Cost\|Monthly Total Expense | cost\|monthly_total_expense |

**Horizon Metrics: Analyzing Price**

| Metric Name Per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Horizon RDS Farm** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type:Horizon RDS Host** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type: Horizon User** | | |
| Total Price | Summary\|Total Price | summary\|vdiTotalPrice |
| **Object Type: Horizon VDI Application Session** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type: Horizon VDI Desktop Session** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |
| **Object Type: Horizon VDI Pool** | | |
| Total Price | Summary\|Total Price | summary\|totalPrice |

**Horizon Metrics : Analyzing Cost**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Horizon RDS Farm** | | |
| Total Cost | Cost \| Monthly Total Cost | farmRDSHostCost\| totalRDSHostCostPerMonth |
| Monthly Projected Total Cost | Cost \| Monthly Projected Cost | farmRDSHostCost\| totalProjectedRDSHostCostPerMonth |

*Table continued on next page*

*Continued from previous page*

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Horizon RDS Host** | | |
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | rdsHostCost\| rdsHostMonthlyEffectiveProjectedTotal Cost |
| MTD Cost | Cost\|MTD Total Cost | rdsHostCost\|rdsHostMTDTotalCost |
| **Object Type: Horizon User** | | |
| Total Cost | Cost\|VDI\|Monthly Total Cost | userCost\|vdiUserCost\| totalUserCostPerMonth |
| Monthly Projected Total Cost | Cost\|VDI\|Monthly Projected Cost | userCost\|vdiUserCost\| totalProjectedUserCostPerMonth |
| **Object Type: Horizon VDI Application Session** | | |
| Monthly Projected Total Cost | Cost\|Monthly Effective Projected Total Cost | vdiApplicationSessionCost\| vdiAppSessionMonthlyEffectiveProject edTotalCost |
| MTD Cost | Cost\|MTD Total Cost | vdiApplicationSessionCost\| vdiAppSessionMTDTotalCost |
| **Object Type: Horizon VDI Desktop Session** | | |
| Monthly Projected Total Cost | Cost\|Monthly Projected Cost | vdiSessionCost\| totalProjectedVDISessionCostPerMont h |
| Total Cost | Cost\|Monthly Total Cost | vdiSessionCost\| totalVDISessionCostPerMonth |
| **Object Type: Horizon VDI Pool** | | |
| Monthly Projected Total Cost | Cost\|Monthly Projected Cost | poolDesktopCost\| totalProjectedDesktopCostPerMonth |
| Total Cost | Cost\|Monthly Total Cost | poolDesktopCost\| totalDesktopCostPerMonth |

**Cloud Director Adapter: Analyzing Price**

| Metric Name per Cost Analysis UI | Metric Name | Metric Key |
|---|---|---|
| **Object Type: Organization** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |
| **Object Type: Organization VDC** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |
| **Object Type: vApp** | | |
| Total Price | Summary\|Metering\|Total Cost | summary\|metering\|value |

### New or Edit Cost Analysis

As part of analyzing cost, New Analysis is the page you use to fill in the details for your cost analysis scenario. You can select a metric, an object type, and the time range for which you want to analyze cost. Cost analysis are rendered using lists or bar graphs.

**Where You Can Create or Edit Cost Analysis**

In the left menu, click **Capacity › Cost**, and then click the **Analysis** tab. The **Cost Analysis** page opens on the right. Click **Add** to create an analysis.

**Table 142: New or Edit Analysis Options**

| Option | Description |
|---|---|
| Analysis Name | Provide a name for your cost, price, or VMware bill analysis. |
| Metrics | Select the metric options based on the type of analysis you want to create: Cost, Price, or VMware Bills analysis.<br><br>**Cost Analysis**<br><br>Select a metric from the list based on which you can analyze and compare costs. Here is the list:<br><br>• Total Cost<br>• Realized Saving<br>• Potential Saving<br>• MTD Cost<br>• Monthly Projected Total Cost<br>For metric name to metric key mapping for each adapter type, see Cost Analysis - Metrics and Metric Keys. |
| | **Price Analysis**<br><br>Select the Total Price metric based on which you can analyze and compare prices. The Total Price metric is based on the pricing policy defined by the user.<br><br>For metric name to metric key mapping for each adapter type, see Cost Analysis - Metrics and Metric Keys. |
| | **VMware Cloud Bills Analysis**<br>Select a metric from the list based on which you can analyze and compare expenses. Here is the list:<br><br>• Total Component Expense<br>• Bill Line Item Expense<br>• Outstanding Expense<br>• Monthly Commit Expense<br>• Monthly On Demand Expense<br>• Monthly Total Expense<br>For metric name to metric key mapping for each adapter type, see Cost Analysis - Metrics and Metric Keys. |
| Object/Group Type | Select an object or group type for which you want to calculate cost. Object types are filtered and displayed based on your metric selection. The objects for which the metric is applicable are displayed. For example, you can |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | find out the Total Cost of a VM by selecting the Total Cost as a metric and VM as an object type.<br>You can also select a Group from the Container drop-down option. |
| Time Range | Select the time range for which you want to analyze and calculate costs. The time range displayed depends on the metric and object type you select.<br>**NOTE**<br>To view the monthly cost breakup for your historical data, set the time range to **3 months**, **6 months**, or **12 months**. |
| Analysis Scope | If you selected an object from **vCenter** in the Object/Group Type field:<br>• Select or enter the specific object type to narrow down your scope of analysis. You can also click the **Select** option to select a specific object type based on which the analysis is created.<br>If you selected a group from **Container** in the Object/Group Type field:<br>• This field is not activated. When you run the analysis, you will see a list of custom groups with the selected type. For example, Department, Function, Location, etc. |
| View Type | Select the view type:<br>• By default, **Top N** is selected to display the top 10 metrics for the selected object.<br>• Select **Comparison** and then select **Metric 2** to run a comparison between the two metrics for the selected object. You can run the comparison analysis for upto 20 objects, click **Add Objects** to add the objects. |
| Run | Click to run the analysis. You see two kinds of results, a list or a bar chart. For more information, see the **Results: List or Compare** section below this table. |
| Save | Click to save the analysis. You can view saved the analysis from the **Saved Analysis** tab. For more information, see Saved Analyses. |

## Results: List or Compare

### Lists

The list results are displayed when you **Run** the Analysis. The list allows you to view the top and bottom 5, 10, and 20 results of your analysis. You can see a list or a bar graph of the selected object type defined in your analysis and the associated metrics for each object type. For example, if you selected Total Cost as a metric and VM as an object type, your list will display the all the VMs and their associated Total Cost.

### Compare

You can also compare costs, prices, or expenses by selecting specific objects from your defined analysis. Click **Select Objects** to select specific objects for comparison. You can see a list or a bar graph of the selected object from your defined analysis and the associated metrics for each object.

**Saved Analyses**

The analyses that you save are listed in the **Saved Analysis** tab.

**Where You Find Saved Analyses**

In the left menu, click **Capacity › Cost**, and then click the **Analysis** tab.

**Saved Analyses Options**

| Option | Description |
|---|---|
| Run | Select and run a single analysis. |
| Horizontal Ellipsis > Delete | Select and delete multiple analyses. |
| Vertical Ellipsis > Run | Select and run the analysis. |
| Vertical Ellipsis > Edit | Select and edit the analysis. For more information, see New or Edit Cost Analysis. |
| Vertical Ellipsis > Delete | Select and delete the analysis. |

## Cost Settings for Financial Accounting Model

You can configure Server Hardware cost driver and resource utilization parameters to calculate the accurate cost and improve the efficiency of your environment.

Cost Drivers analyze the resources and the performance of your virtual environment. Based on the values you define, Cost Drivers can identify reclamation opportunities and can provide recommendations to reduce wastage of resources and cost.

**How to Set Your Depreciation Model and Years for vCenter**

You can set your depreciation model and years using the following steps.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. Click **Settings**.
3. Select vCenter as **Infrastructure Type** from the drop-down menu.
4. In **Cost Settings - Financial Accounting Model** page, select the **Depreciation Years** between two and five.
5. Select the **Depreciation Model** as per your requirement and click **Save**.

**Editing Cost Ratio for VMware Cloud on AWS, Azure VMware Solution, and Google Cloud VMware Engine**

You can set or modify the cost ratio for public cloud account using the following steps.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. Click **Settings**.
3. Select the public cloud account of your choice from the **Infrastructure Type** drop-down menu.
4. Select the required organization from the **Organization** drop-down menu.
5. Enter the cost ratio for **CPU**, **Memory**, and **Storage**.
6. Enter the Discount percentage.

**NOTE**
Discount is applied only when billing is deactivated on VMC on AWS/ Azure VMware Solution/ Google Cloud VMware Engine adapter and it is applied on the reference cost (list price) per hour applicable for that region.

7. Select the **Region** and click **Save**.

**NOTE**
This value is considered when VMware Cloud vCenters are configured to VMware Aria OperationsVMware Cloud Foundation Operations using vCenter adapters only.

The default discount % value is zero. You can set or edit the discount % for all the organizations in any of your VMware Cloud environments or you can set or edit the discount % for specific organization in any of your VMware Cloud environments. You can run cost calculation and check whether the discount % is reflected in the Monthly CPU Base Rate, Monthly Memory Base Rate, and Monthly Storage Base Rate metrics.

## Configuring Depreciation Preferences

To compute the amortized cost of the Server Hardware cost driver, you can configure the depreciation method and the depreciation period. Cost Drivers supports two yearly depreciation methods and you can set the depreciation period from two to five years.

**NOTE**
Cost Drivers calculates the yearly depreciation values and then divides the value by 12 to arrive at the monthly depreciation.

| Method | Calculation |
|---|---|
| Straight line | `Yearly straight line depreciation = [(original cost - accumulated depreciation) / number of remaining depreciation years]` |
| Max of Double or Straight | `Yearly max of Double or Straight = Maximum (yearly depreciation of double declining balance method, yearly depreciation of straight line method)`<br>`Yearly depreciation of double declining method= [(original cost - accumulated depreciation) * depreciation rate].`<br><br>`Depreciation rate = 2 / number of depreciation years.`<br>**NOTE**<br>`Double declining depreciation for the last year = original cost - accumulated depreciation` |

## Example for Straight Line Depreciation Method

| Year | Original Cost | Accumulated Depreciation | Straight Line Depreciation Cost |
|---|---|---|---|
| Year 1 | 10000 | 0 | `[(10000-0)/5] = 2000` |
| Year 2 | 10000 | 2000 | `[(10000-2000)/4] = 2000` |
| Year 3 | 10000 | 4000 | `[(10000-2000)/3] = 2000` |
| Year 4 | 10000 | 6000 | `[(10000-2000)/2] = 2000` |
| Year 5 | 10000 | 8000 | `[(10000-2000)/1] = 2000` |

## Example for Max of Double and Straight Line Depreciation Method

| Year | Original Cost | Depreciation Rate | Accumulated Depreciation | Straight Line Depreciation Cost |
|---|---|---|---|---|
| Year 1 | 10000 | 0.4 | 0 | `Maximum([(10000-0)*0.4], [(10000-0)/5])`<br>`= Maximum(4000, 2000) = 4000`<br>which is 333.33 per month. |
| Year 2 | 10000 | 0.4 | 4000 | `Maximum([(10000-4000)*0.4], [(10000-4000)/4])`<br>`= Maximum (2400, 1500) = 2400`<br>which is 200 per month. |
| Year 3 | 10000 | 0.4 | 6400 | `Maximum([(10000-6400)*0.4], [(10000-6400)/3])`<br>`= Maximum (1440, 1200) = 1440`<br>which is 120 per month. |
| Year 4 | 10000 | 0.4 | 7840 | `Maximum([(10000-7840)*0.4], [(10000-7840)/2])` |

*Table continued on next page*

*Continued from previous page*

| Year | Original Cost | Depreciation Rate | Accumulated Depreciation | Straight Line Depreciation Cost |
|---|---|---|---|---|
| | | | | `= Maximum (864, 1080) = 1080`<br>which is 90 per month. |
| Year 5 | 10000 | 0.4 | 8920 | `Maximum([(10000-8920)*0.4], [(10000-8920)/1])`<br>`= Maximum (432, 1080) = 1080`<br>which is 90 per month. |

## Overview of Cost Drivers

Cost Drivers are the aspect that contributes to the expense of your business operations. Cost drivers provide a link between a pool of costs. To provide a granular cost visibility and to track your expenses of virtual machines accurately in a private cloud, VMware Aria OperationsVMware Cloud Foundation Operations has identified eight key cost drivers. You can see the total projected expense on your private cloud accounts for the current month and the trend of cost over time.

You can now set a total cost for the License, Labor, Network, Maintenance, and facilities cost drivers in VMware Aria OperationsVMware Cloud Foundation Operations:

> **NOTE**
> The total cost set by you is distributed across resources in the data center. For example, if you set the total cost for the RHEL license, the cost is divided across all the hosts and VMs which use the RHEL license.

According to the industry standard, VMware Aria OperationsVMware Cloud Foundation Operations maintains a reference cost for these cost drivers. This reference cost helps you for calculating the cost of your setup, but might not be accurate. For example, you might have received some special discounts during a bulk purchase or you might have an ELA with VMware that might not match the socket-based pricing available in the reference database. To get accurate values, you can modify the reference cost of cost drivers in VMware Aria OperationsVMware Cloud Foundation Operations, which overrides the values in the reference database. Based on your inputs, VMware Aria OperationsVMware Cloud Foundation Operations recalculates the total amount for the private cloud expenses. After you add a private cloud into VMware Aria OperationsVMware Cloud Foundation Operations, VMware Aria OperationsVMware Cloud Foundation Operations automatically discovers one or more vCenter Servers that are part of your Private Cloud. In addition, it also retrieves the inventory details from each vCenter Server. The details include:

- Associated clusters: Count and names
- ESXi hosts: Count, model, configuration, and so on.
- Data stores: Count, storage, type, capacity
- VMs: Count, OS type, tags, configuration, utilization

Based on these configuration and utilizations of inventory, and the available reference cost, VMware Aria OperationsVMware Cloud Foundation Operations calculates the estimated monthly cost of each cost driver. The total cost of your private cloud is the sum of all these cost driver expenses.

You can modify the expense of your data center. These costs can be in terms of the percentage value or unit rate, and might not always be in terms of the overall cost. Based on your inputs, the final amount of expense is calculated. If you do not provide inputs regarding expenses, the default values are taken from the reference database.

To know more about Reference Cost, see Reference Based Costing for / .

You can see the projected cost of private cloud for the current month and the trend of total cost over time. For all the expenses, cost drivers in VMware Aria OperationsVMware Cloud Foundation Operations display the monthly trend of the cost variations, the actual expense, and a chart that represents the actual expense and the reference cost of the expense.

> **NOTE**
> If the vCenter was added from more than six months, the trend displays the total cost for the last six months only. Otherwise, the trend displays the total cost from the month the vCenter was added into VMware Aria OperationsVMware Cloud Foundation Operations.

Infrastructure Type

You have the option to select the infrastructure type as either vCenter or VMC on AWS, based on your selection the cost drivers are displayed on the Cost Drivers page. You can add or edit the cost drivers as per your requirement.

> **NOTE**
> You can edit the cost driver values either in All Datacenter mode or Specific Datacenter mode. Ensure that you download and upload the cost driver configuration file in the same mode (either All DC mode or Specific DC mode).

For the vCenter infrastructure type the following private cloud cost drivers are applicable.

- Server Hardware : Traditional
- Server Hardware : Hyper - Converged
- Storage
- License
- Applications
- Maintenance
- Labor
- Network
- Facilities
- Additional Cost

For the VMC on AWS infrastructure type the following private cloud cost drivers are applicable.

- License
- Additional Cost

All other costs from VMC that are not directly attributed to specific hosts like load balancer, Tax, and other costs are grouped under additional cost driver and equally distributed among all hosts.

Export and Import Cost Drivers

Except for additional cost drivers, you can export or import the remaining cost drivers associated with your private cloud. With this functionality you can edit the cost driver values from the excel sheet instead of editing them from the user interface. You have an option to select all cost drivers and export them or you can select individual cost drivers and export them.

> **NOTE**
> The import and export functionality is applicable only to vCenter cost drivers, the functionality is not available for VMware Cloud on Amazon Web Services.

**Table 143: Expense Types**

| Cost Drivers | Description |
|---|---|
| Select Datacenter | The Select Datacenter option allows you to choose the data center for which the cost driver changes are applicable.<br>**NOTE**<br>You can select a specific data center and modify the cost driver values of that data center, or you can modify the cost drivers and apply the changes to all the data centers. |
| Export | Click export to export the cost details for all the cost drivers. You can select individual cost drivers and export them also.<br>For more information, see Import or Export Cost Drivers. |
| Import | Click import to browse and upload the updated cost driver configuration file (xls/csv). The import cost driver file should have the same template as that of the exported file.<br>**NOTE**<br>You might be prompted with error messages if the uploaded file has errors. You can ignore the error or you can download the log file. You can click ignore error to omit the incorrect values and include the correct ones.<br>For more information, see Import or Export Cost Drivers. |
| **Server Hardware : Traditional** | The Server Hardware cost driver tracks all the expenses for purchasing of hardware servers that are part of vCenter Servers. You see the server cost based on CPU age and server cost details.<br>**NOTE**<br>You can now select an individual server from the server group and specify the unique cost for each individual server.<br>For more information, see Editing Server Hardware : Traditional. |
| **Server Hardware : Hyper-Converged** | The Server Hardware : Hyper-Converged cost driver, tracks the expenses associated with hyper converged infrastructure components. The Server Hardware : Hyper-Converged cost driver includes expenses for the Hyper Converged servers like vSAN activated servers and vXRail. The expense provided is for both compute and storage.<br>**NOTE**<br>The customizations that were performed for vSAN server costing under Server Hardware : Traditional in the earlier versions will not be carried forward to 7.5 as the vSAN activated servers will fall under Server Hardware : Hyper-Converged servers now.<br>For more information, see Editing Server Hardware: Hyper-Converged. |
| **Storage** | You can calculate the storage cost at the level of a data store based on the tag category information collected from vCenter Server. You see the storage total distribution based on category and the uncategorised cost details.<br>**NOTE**<br>The vSAN data stores are not displayed as part of this cost driver page.<br>For more information, see Edit Monthly Cost of Storage. |
| **License** | You see the licenses cost distribution for the operating systems cost and VMware license of your cloud environment.<br>**NOTE**<br>For Non-ESX physical servers, VMware license is not applicable. |

*Table continued on next page*

*Continued from previous page*

| Cost Drivers | Description |
|---|---|
| | For more information, see Edit Monthly Cost of License. |
| **Maintenance** | You see the maintenance cost distribution for the server hardware and operating system maintenance. You can track your total expense with hardware and operating system vendors. For more information, see Edit Monthly Cost of Maintenance. |
| **Labor** | You see the labor cost distribution for the servers, virtual infrastructure, and operating systems. You can view the total administrative cost for managing physical servers, operating systems and virtual machines. You can track all expenses spent on human resources to manage the data centers.<br>**NOTE**<br>• Labor cost includes expenses on backup appliance virtual machine (VDP virtual appliance).<br>• For physical servers, operating system labor cost and servers labor costs are applicable, virtual infrastructure cost is not considered.<br>For more information, see Edit Monthly Cost of Labor. |
| **Network** | You see the networks costs by NIC type. You can track a network expense based on different types of NICs attached to the ESX server. You can view the total cost of physical network infrastructure that includes the internet bandwidth, and is estimated by count and type of network ports on the ESXi Servers.<br>**NOTE**<br>For physical servers, the network details are not captured. So, the network cost is considered as zero.<br>For more information, see Edit Monthly Cost of the Network. |
| **Facilities** | You see the cost distribution for the facilities such as real estate costs, such as rent or cost of data center buildings, power, cooling, racks, and associated facility management labor cost. You can point to the chart to see the cost details for each facility type.<br>For more information, see Edit Monthly Cost of Facilities. |
| **Additional Cost** | You can see the additional expenses such as backup and restore, high availability, management, licensing, VMware software licensing.<br>For more information, see Editing Additional Costs. |
| **Application Cost** | You can see the cost of different application services you are running in your environment compared to your overall expenses. Some examples of application cost are, cost of running SQL server cluster and cost of running Anti-virus on VMs.<br>For more information, see Edit Application Cost. |

You can select a data center to view the information specific to the data center.

### Import or Export Cost Drivers

The cost driver editing process has been enhanced to support export and import of existing cost driver configurations. You can download (export) the existing cost driver configurations as an xls/csv file, edit the cost drivers and import the updated file back to the system. You must ensure that the import cost driver file should have the same template as that of the exported file.

**Where to find the Import or Export Option**

From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**. In the **Cost Drivers** tab, select either **Import** or **Export**.

The import and export functionality is applicable only for vCenter cost drivers, the functionality is not available for VMware Cloud on Amazon Web Services.

Using the import and export option, you can perform the following actions on the cost drivers:

- Export or import the cost driver configuration file.
- Read and edit the cost driver configuration file.
- Validate the updated cost driver configuration file and report errors.
- Identify the error from the log file and correct the errors.

You are prompted with error messages if the uploaded file has errors. You can correct the errors and upload the file, or you can ignore the errors, the system still allows you to upload the file.

## Cloud Providers Overview

By default, you can see that Google Cloud, IBM Cloud, VMware Cloud on AWS, Azure VMware Solution, and Google Cloud VMware Engine are included in VMware Aria OperationsVMware Cloud Foundation Operations. You can also add your own cloud provider by using a standard VMware Aria OperationsVMware Cloud Foundation Operations template.

You can configure the new cloud provider as per the standard VMware Aria OperationsVMware Cloud Foundation Operations template and perform a migration scenario. The VMware Aria OperationsVMware Cloud Foundation Operations template contains data points for vCPU, CPU, RAM, OS, region, plan term, location, and built-in instance storage, you must provide these values when you add cloud providers. The result of the migration scenario helps you assess the cost savings achieved using your cloud provider against the default cloud providers.

You can edit the rate card for new cloud providers and default cloud providers. However, you cannot delete the default cloud providers.

**Add or Edit Cloud Provider**

You can use the Add Cloud Provider workspace to add or edit a cloud provider. You can edit the cloud provider rate card for default cloud providers and add the new cloud provider.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers › Cloud Providers**.
2. To add or edit the cloud provider, click **Add** icon or **Edit** option from the vertical ellipsis menu.
3. Enter or edit the **Cloud Provider Name**.
4. Select the cloud provider logo and click **Upload Logo**.
5. Click **Next**.
6. Click **Download Template** to edit the **Instance Configuration** and **Pricing Configuration** values.

   **NOTE**
   When you edit a cloud provider the Download Template link is replaced with Download Existing Rate Card. You can add, edit, or delete the rows with the instance names. You can fill in the instance settings values and the mandatory price configuration values as mentioned in the template.

7. Select the updated template and click **Upload Rate Card**.
8. Click **Validate**.

   **NOTE**
   VMware Aria OperationsVMware Cloud Foundation Operations validates the rate card and reports success or failure. If errors are reported, you can correct the errors and proceed further.

9. Click **Finish**.

The new cloud provider is now part of the VMware Aria OperationsVMware Cloud Foundation Operations cloud provider list.

## Billing Enhancements for Horizon Management Pack and Virtual Hosts

The cost calculation of VMware Aria OperationsVMware Cloud Foundation Operations has been enhanced to include the end point objects of Horizon Management Pack and virtual hosts. Earlier, the cost calculation was based on the metrics collected for each end point object.

The cost calculation for the end point objects is now based on the following criteria:

- Each Virtual Desktop Infrastructure Virtual Machine ( (VDI VM)) is counted as 0.25 Operating System Instance (OSI)
- Each Remote Desktop Service Host (RDS Host) is counted as 0.25 Operating System Instance
- One Operating System Instance for each Connection Server
- Virtual Hosts (ESXi hosted on a VM) is not counted against license usage
- VMs hosting the virtual hosts are counted against license usage

There are no VDI VM objects discovered by Horizon MP. VMware Aria OperationsVMware Cloud Foundation Operations reports the number of VDI VMs in the bill. The number of VDI VMs appear under Virtual Machine node of the vCenter MP.

A new property `Is Horizon Managed` is introduced which is published by vCenter Adapter on VMs, Hosts and Clusters. VDI VMs are identified by this property (if the property is true for the VM, then the VM is considered as a VDI VM). If the property is true at VM level, then the same property is published for the cluster and all the hosts under that cluster, indicating that those hosts and clusters are Horizon Managed objects.

> **NOTE**
> VMs and Hosts are considered as Horizon Objects only when Horizon Management pack is installed. If Horizon Management pack is not installed the objects are considered as Non-Horizon Objects irrespective of the object properties published by vCenter Adapter.

**How to Identify the Virtual Host**

You can identify the virtual hosts by the following property.

- Hardware |Vendor = "VMware, Inc"

## Editing Cost Drivers

You can manually edit monthly cost of all the eight expense types from the current month onwards.

The configuration used for cost drivers determines how VMware Aria OperationsVMware Cloud Foundation Operations calculates and displays the cost.

**Editing Server Hardware : Traditional**

You can view, add, edit, or delete the cost of each server group, based on their configuration and the purchase date of a batch server running in your cloud environment. You can also specify the server cost for individual servers in a server group. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Drivers tab, click **Server Hardware : Traditional**.

**NOTE**
You can customize the default value of cost per server and specify exclusive values for other servers in the list.

For example, if you have a system that has eight servers you can modify the default reference value from $1000 to $800 for eight servers. You can also select two servers from the list and customize their value as $600. So, any new server that is added to the system will have the default value as $800.

3. Select the required edit mode for changing the server hardware cost.
   - **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   - **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

     **NOTE**
     When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4. Click any server from the list of **Server Group Description**.
   The cost drivers groups all server hardware from all data centers in your inventory based on their hardware configuration.

| Category | Description |
|---|---|
| Server Group Description | Displays the name of the server in your inventory. |
| Number of Servers | Displays the total number of servers of any particular hardware configuration in your inventory. |
| Monthly Cost | Displays the average monthly cost for server. This value is calculated as a weighted average of prices of purchased and leased batches. |

   **NOTE**
   If the vCenter does not set the server vendor model, then that server is listed under the group Others in VMware Aria OperationsVMware Cloud Foundation Operations for **Server Hardware : Traditional** cost driver and **Server Hardware : Hyperconverged** cost driver.

5. After selecting a server group, you can manually enter the required fields.
   a) Enter the Purchase Type and Cost Per Server.

      **NOTE**
      You can use the **+ ADD COST PER SERVER** option to create multiple server batches and set the cost for a specific server in a server group.

   b) Click **Save**.

### Editing Server Hardware: Hyper-Converged

You can view, add, edit, or delete the cost of Hyper converged Infrastructure (HCI) component in your server group. You can specify the cost per server and compute percentage exclusively for the HCI servers. After you update the server hardware cost, cost drivers update the total monthly cost and average monthly cost for each server group.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Drivers tab, click **Server Hardware : Hyper-Converged**.
3. Select the required edit mode for changing the server hardware cost.

- **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
- **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

  **NOTE**
  When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4. Click any server from the list of **Server Group Description**.
   The cost drivers groups all server hardware from all data centers in your inventory based on their hardware configuration.

   **NOTE**
   If the vCenter does not set the server vendor model, then that server is listed under the group Others in VMware Aria OperationsVMware Cloud Foundation Operations for **Server Hardware : Traditional** cost driver and **Server Hardware : Hyperconverged** cost driver.

| Category | Description |
|---|---|
| Server Group Description | Displays the name of servers falling under vSAN clusters and vXrail servers in your inventory. |
| Number of Servers | Displays the total number of servers of any particular hardware configuration in your inventory. |
| Monthly Cost | Displays the average monthly cost for server. This value is calculated as a weighted average of prices of purchased and leased batches. |

   **NOTE**
   You can edit the Compute Pct column to adjust the storage rate of the vSAN datastores. You can use the same percentage to determine the cost.

5. After selecting a server group, you can manually enter the required fields.
   a) Enter Purchase Type, Cost Per Server, and Compute Percentage.

      **NOTE**
      You can use the **+ ADD COST PER SERVER** option to create multiple server batches and to customize the cost per server.

   b) Click **Save**.

### Edit Monthly Cost of Storage

The storage hardware is categorized according to the datastore tag category. You can edit the monthly cost per storage GB for the datastores based on their storage category (using tags) and storage type (NAS, SAN, Fiber Channel, or Block).

To edit the cost based on the storage category, you must create tags and apply them to the datastores on the vCenter user interface. For more information, see the VMware vSphere Documentation.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Drivers tab, click **Storage**.
3. Select a tag category.
   Assume that you have two tag categories (for example, Profile and Tiers) with three tags in each category, you can select either Profile or Tiers from **Tag Category** to categorize the datastores based on tags.

| Category | Description |
|---|---|
| Edit Mode | You can select the storage cost to be applicable for all the data centers or a specific data center.<br>• **Edit for All Data Centers** mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.<br>• **Edit for specific Data Center** mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost. |
| Select Data center | You can select the data center for which you want to change the storage cost. This field is applicable only for specific data centers. |
| Tag Category | • **Category** displays the tag categories for datastores and also the tags associated with the category. |
| Datastores | Displays the total number of datastores for a specific category or type. You can click the datastore value to see the list of datastores and its details such as monthly cost, total GB for each datastore. |
| Total Storage (GB) | Displays the total storage for a specific category or type. |
| Monthly Cost Per GB | Displays the monthly cost per GB for a specific category or type. You can edit this value for defining the monthly cost per GB for datastores. |
| Monthly Cost | Displays the total monthly cost for a specific category or type. |

4. Click **Save**.

### Edit Monthly Cost of License

You can edit the different license cost of your VMware Aria OperationsVMware Cloud Foundation Operations environment. You can now set a total fixed cost for the license in VMware Aria OperationsVMware Cloud Foundation Operations. The total license cost is divided across all the hosts present in the data center. You can edit the license cost by either selecting the ELA charging policy or selecting the per socket value.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Drivers tab, click **License**.
3. Select the required edit mode for changing the license cost.
   – **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   – **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

     **NOTE**
     When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4. Click **Save**.
   The Cost drivers display all the licenses in your VMware Aria OperationsVMware Cloud Foundation Operations environment.

| Category | Description |
| --- | --- |
| Name | Displays the different license categories:<br>• VMware by Broadcom Licenses.<br>– **VMware Software Per Core**: The VMware Software Per Core license is applicable for ESXi licensing and covers the monthly cost of the VMware Cloud Foundation (VCF) license for VMware ESXi server 8.0 and later is charged by Per Core. The VCF license is assumed by default for all core-based licenses after you upgrade to VMware ESXi server 8.0.<br>**NOTE**<br>The cost value of $9.7 has been calculated based on the three-year term license pricing value ((350/3)/12). If you are using a vSphere Foundation (VVF) license, you must update this value to $3.75 ((135/3)/12) .<br>– **VMware Software**: The VMware Software license is applicable for ESXi licensing only. The monthly cost of VMware Software for VMware ESXi server 7.0 or earlier is charged by Per Socket. The cost value of $51.32 has been calculated using the standard license pricing value.<br>– **VMware vSAN**: vSAN license is applicable only to ESXi servers that are classified as Server Hardware- Hyper Converged and this cost is distributed to all the datastores under these servers. There are two cost components, Monthly cost of VMware vSAN Per Socket and Monthly cost of VMware vSAN SnS have been included for the vSAN cost calculation. The default values for these components are based on the reference database values.<br>**NOTE**<br>VMware vSAN licenses are included in the VCF licenses. If you are using a VCF license, update the cost value as applicable to avoid adding extra cost.<br>• **Windows Server**: The licensing cost for the Windows operating system falls under one of the following categories:<br>Per Core License, applicable for<br>• Windows Server 2016<br>• Windows Server 2019<br>Per Socket License, applicable for<br>• Windows NT 4.0<br>• Windows Server 2003 |

*Table continued on next page*

*Continued from previous page*

| Category | Description |
|---|---|
| | • Windows Server 2008<br>• Windows Server 2012<br><br>Per Instance License, applicable for<br>• Windows XP<br>• Windows Vista<br>• Windows 98<br>• Windows 95<br>• Windows 8<br>• Windows 7<br>• Windows 3.1<br>• Windows 2000<br>• Windows 10<br>• **SUSE** or **Red Hat**: Cost drivers categorize the Linux operating system under **SUSE** or **Red Hat**.<br>• **Other Operating Systems**: Displays the category of the operating system. If the operating system is not Windows or Linux, cost drivers categorize the operating system under **Other Operating Systems**. |
| VMs | Displays the number of virtual machines that are running on the specific operating system. |
| Sockets | Displays the number of sockets on which the specific operating system is running. |
| Charged by | Displays whether a cost is charged by socket or ELA.<br><br>**NOTE**<br>The Charged By column can be edited to mention that the cost is charged by socket, core, instance, or ELA. |
| Total Cost | Displays the total cost of the specific operating system. |

5. Click **Save**.

According to your inputs, VMware Aria OperationsVMware Cloud Foundation Operations calculates and displays the total cost and updates the Charged by column with the option that you have selected.

**Customizing License Assignment**

You can customize the licensing cost associated with your host using the custom license assignment option. Based on your requirement you can add or delete different operating system licenses to your host. With the custom license assignment option, you can increase or decrease the licensing cost associated with your host.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Drivers tab, click **License**.
3. Select the required edit mode for changing the monthly license cost.
   – **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   – **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

**NOTE**
When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4. To customize the license cost for a specific server, click **Customize License Assignment**.
5. Select the host for which you want to customize the license cost and click **Assign**.
6. From the drop-down menu, select the operation system and click **OK**.

   The new operating system is listed under the Current Assignment column.

7. To remove an existing operating system from the host, under **Current Assignment** click X icon next to the operating system.

   The license cost of the removed operating system is reduced from the total cost.

8. Click **Save**.
9. Navigate to the **Cost Calculation Status** tab and click **Run**.

The license cost is updated for the host, the * sign next to the host indicates that the license cost for the host has changed.

| Category | Description |
|---|---|
| Server | You can select the server for which you want to customize the license cost. |
| Current Assignment | Displays the current operating systems associated with the host. |
| Default Assignment | Displays the default operating systems associated with the host. |
| Filter | Filters the hosts based on the operating system type. |
| Reset | Resets the license cost of the host to the default value. |

### Edit Monthly Cost of Maintenance

You can edit the monthly cost of maintaining your cloud environment. Maintenance cost is categorized into hardware maintenance cost and operating system maintenance cost. Hardware maintenance cost is calculated as a percentage of the purchase cost of servers. Operating system maintenance cost is calculated as a percentage of the Windows licensing costs. You can snow specify a total fixed cost for maintenance in VMware Aria OperationsVMware Cloud Foundation Operations. The total maintenance cost is divided across all the hosts present in the data center.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Drivers tab, click **Maintenance**.
3. Select the required edit mode for changing the monthly maintenance cost.
   - **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   - **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

     **NOTE**
     When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4. To customize the maintenance cost for a specific server, click **Edit For Individual Servers**.
5. Click **+Add Cost Per Server**.

6. From the **Select Server's for customization** drop-down select the required server and click **OK**.
7. Specify the Server Hardware Percentage and OS Percentage and click **Save**.
   View the change in maintenance cost after you have run the cost calculation cycle.

### Edit Monthly Cost of Labor

You can edit the monthly cost of labor for your cloud environment. You can set a total fixed cost for labor in VMware Aria OperationsVMware Cloud Foundation Operations. The total labor cost is divided across all the hosts present in the data center. The labor cost is combination of the total cost of the server administrator, virtual infrastructure administrator, and the operating system administrator.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Driver tab, click **Labor**.
3. Select the required edit mode for changing the monthly labor cost.
   – **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   – **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

   **NOTE**
   When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu

4. Edit the monthly labor cost.
   • Edit the detailed cost of labor.
   • Edit the total monthly labor cost for servers, virtual infrastructure, and operating system.

5. To customize the labor cost for a specific server, click **Server** and then click **Edit For Individual Servers**.
6. Click **+Add Cost Per Server**.
7. From the **Select Server's for customization** drop-down select the required server and click **OK**.
8. Specify the Monthly hours of labor per hour, Labor hourly rate, and click **Save**.
   The monthly labor cost is displayed.

| Category | Description |
|---|---|
| Category | Displays the categories of labor cost, servers, virtual infrastructure, and operating system |
| Calculated by | Displays whether the cost is calculated hourly or monthly. |
| Total Monthly Cost | Displays the total monthly cost of the particular category |
| Reference Cost | Displays the reference cost for the category from the cost drivers database |

The total monthly cost is updated. The hourly rate option or the monthly cost option that you select is updated in the **Calculated by** column.

### Edit Monthly Cost of the Network

You can edit the monthly cost for each Network Interface Controller (NIC) type or can edit the total cost of all the networking expenses associated with the cloud. You can now set a total fixed cost for network resources in VMware Aria OperationsVMware Cloud Foundation Operations. The total network cost is divided across all the hosts present in the data center.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.

2. In the Cost Driver tab, click **Network**.
3. Select the required edit mode for changing the monthly network cost.
   – **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   – **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.

     **NOTE**
     When you select Edit for specific data center as the edit mode, then the select data center option is activated. Select the data center from the drop-down menu
4. Edit the monthly cost of network.
   – Modify the values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and the 100 Gigabit NIC.
   – Modify the total monthly cost of all network expenses associated with the cloud.
5. To customize the network cost for a specific server, click **Edit For Individual Servers**.
6. Click **+Add Cost Per Server**.
7. From the **Select Server's for customization** drop-down select the required server and click **OK**.
8. Specify values for 1 Gigabit NIC, 10 Gigabit NIC, 25 Gigabit NIC, 40 Gigabit NIC, and 100 Gigabit NIC and click **Save**.
   View the change in network cost after you have run the cost calculation cycle.

### Edit Monthly Cost of Facilities

For your cloud environment, you can specify the total monthly cost of facilities or edit the facilities cost for real estate, power, and cooling requirements. You can now set the total fixed cost for facilities in VMware Aria OperationsVMware Cloud Foundation Operations. The total facilities cost is divided across all the hosts present in the data center.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Driver tab, click **Facilities**.
3. Select the required edit mode for changing the monthly facilities cost.
   – **Edit for All Data centers** - mode helps you to customize a single cost driver value for all the data centers. Any customizations done for the Specific data center mode are lost.
   – **Edit for specific Data Center** - mode helps you to customize different cost driver values for different data centers. Any customizations done for All data centers mode are lost.
4. Select the data center from the drop-down menu.

   **NOTE**
   If you select Edit for specific data center as the edit mode, then the select data center option is activated.

5. Edit the monthly facilities cost.
   – Modify the cost of rent or real estate per rack unit and modify the monthly cost of power and cooling per kilowatt-hour.
   – Modify the total monthly cost of facilities.
6. To customize the facilities cost for a specific server, click **Edit For Individual Servers**.
7. Click **+Add Cost Per Server**.
8. From the **Select Server's for customization** drop-down select the required server and click **Ok**.
9. Specify the Cost Per Kilowatt and Real Estate Cost Per Rack Unit and click **Save**.
   View the change in network cost after you have run the cost calculation cycle.

### Editing Additional Costs

The additional cost lets you add any additional or extra expense that is not covered by other expenses categorized by VMware Aria OperationsVMware Cloud Foundation Operations. No reference value is present for this expense.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Driver tab, click **Additional Costs**.
3. Enter or select the cost type for the expenses.

   **NOTE**
   Additional cost driver allows you to assign costs at Host, vCenter, VM, cluster, or data center level. For example, if you want to keep a cluster protected using the disaster recovery services, which involves an additional cost of $5000, you can do that by editing the additional cost driver.

4. Select the **Entity Type** and **Entity Selection**.

   The **Entity Count** gets updated after the cost calculation runs.

   **NOTE**
   For the Cluster, Data center, Host and vCenter entity types, you must enter the the ID entity value. For the VM entity type, you can select one of the following entity values: Custom Properties, Tag, or Guest OS. The Custom Properties are derived from VMware Cloud Foundation Automation, Tags is derived from vSphere, and Guest OS is derived from the VMware tools.

5. Enter the **Monthly Cost per entity** .

   The **Total Cost per month** gets computed automatically.

6. Click **Save**.

   **NOTE**
   After you update the Additional Cost configuration, you must reload the page manually to view the updated values.

#### Edit Application Cost

VMware Aria OperationsVMware Cloud Foundation Operations allows you to edit the application cost of an application present in your cloud environment. You can only modify the cost associated with the application, as all the other attributes are predefined.

Create applications in VMware Aria OperationsVMware Cloud Foundation Operations.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers**.
2. In the Cost Drivers tab, click **Applications**.
3. Click the edit icon next to the application cost you want to edit.

   **NOTE**
   You can now specify the cost of packaged applications that are discovered by the Service Discovery Management Pack. Earlier the option to specify the application cost was available only for business applications defined by the user.

4. Modify the cost of the application.
5. Click **Save**.

## Cluster Cost Overview

VMware Aria OperationsVMware Cloud Foundation Operations calculates the base rates of CPU and memory so that they can be used for the virtual machine cost computation. Base rates are determined for each cluster, which are

homogeneous provisioning groups. As a result, base rates might change across clusters, but are the same within a cluster.

1. VMware Aria OperationsVMware Cloud Foundation Operations first arrives at the fully loaded cost of the cluster from the cost drivers. After the cost of a cluster is determined, this cost is split into CPU and memory costs based on the industry standard cost ratios for the different models of the server.
2. The CPU base rate is first computed by dividing the CPU cost of the cluster by the CPU capacity of the cluster. CPU base rate is then prorated by dividing the CPU base rate by expected CPU use percentage to arrive at a true base rate for charging the virtual machines.
3. The memory base rate is first computed by dividing the memory cost of the cluster by the memory capacity of the cluster. Memory base rate is then prorated by dividing the memory base rate by expected memory use percentage to arrive at true base rate for charging the virtual machines.
4. You can either provide the expected CPU and memory use or you can use the actual CPU and memory usage values.
5. The base rates are monthly rates.

| Cluster Cost Elements | Calculation |
|---|---|
| Total Compute Cost | Total Compute Cost = (Total Infrastructure cost, which is a sum of all cost drivers) – (Storage cost) – (Direct VM cost, which is sum of OS labor, VM labor and any Windows Desktop licenses). |
| Expected CPU and Memory use | Expected CPU and Memory use = These percentages are arrived based on historical actual use of clusters. |
| Per GHz CPU base rate | Per GHz CPU base rate = (Cost attributed to CPU out of Total compute cost) / (Expected CPU Utilization * Cluster CPU Capacity in gHZ). |
| Per GB RAM base rate | Per GB RAM base rate = (Cost attributed to RAM out of Total compute cost) / (Expected Memory Utilization * Cluster RAM Capacity in GB). |
| Average CPU Utilization | Average CPU Utilization = (Cost attributed to CPU utilization of VMs in a cluster, out of Total compute cost) / (Total number of VMs in the cluster). |
| Average Memory Utilization | Average Memory Utilization = (Cost attributed to Memory utilization of VMs in a cluster, out of Total compute cost) / (Total number of VMs in the cluster). |
| Expected CPU Utilization | The utilization percentage level of CPU that the cluster is expected to operate.<br>**NOTE**<br>When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value. |
| Expected Memory Utilization | The utilization percentage level of Memory that the cluster is expected to operate.<br>**NOTE**<br>When you select actual utilization as the cost calculation mode, the cost engine by default rounds off the actual utilization value in multiples of five or to the nearest value. |

### Cluster Cost Computation with Allocation Model

You can now use the allocation model to compute the cost of clusters in VMware Aria OperationsVMware Cloud Foundation Operations, earlier the cluster cost computation was based on the cluster utilization. When you perform cost computation using the allocation model, you can set the over commit ratio for CPU, RAM, and storage.

> **NOTE**
> The allocation ratio can be set at both cluster level and datastore cluster level. You can also mention the storage base rate, which will displayed at the datastore level.

**Table 144: Cluster Base Rate Computation with Allocation Model**

| Base Rate | Formula |
|---|---|
| vCPU Base Rate | vCPU base rate = B1 = (Cost attributed to CPU) / (Number of vCPUs in a cluster) |
| RAM Base Rate | RAM base rate = B2 = (Cost attributed to RAM) / Number of vRAMs in a cluster)<br>**NOTE**<br>The cost computation is based on Over Commit ratio. If the Over Commit ratio is 1:4, and total cores in cluster are 6, then vCPU count = 24, in case if the allocated vCPU exceeds this targeted number, then the maximum value is selected. |

**Table 145: Virtual Machine Cost Computation with Allocation Model**

| Cost | Formula |
|---|---|
| Virtual Machine Cost | Virtual machine cost = (Number of vCPU allocated x B1 of cluster it belongs to) + Number of vRAMs allocated x B2 of cluster it belongs to) + storage cost + direct cost.<br>**NOTE**<br>Storage allocated represents the Storage Base Rate based on allocation. |

### Editing Cluster Cost Calculation Methods

You can edit the cluster cost calculation method based on your business requirement. The cost of a cluster is derived from cost drivers. Virtual machine cost is calculated by multiplying base rates with the utilization of the VMs.

1. From the left menu, click **Operations › Configurations**, and then click **Cost Drivers › Cluster Cost**.
2. Click **CHANGE**.
   The Cluster Cost Calculation Methods dialog box is displayed.
3. Select any one of the Cluster Cost Calculation methods.

| Option | Description |
|---|---|
| Cluster Usable Capacity After HA and Buffer | The cluster cost calculated total capacity minus resources needed for High Availability (HA) and the capacity buffer setting. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Base rates are calculated based on the total cost of the cluster and Usable Capacity after HA and Buffer. Virtual machine costs are calculated from these base rates.<br>Things to note:<br>• A lower buffer reduces the base rates and causes the virtual machines to become cheaper.<br>• A higher buffer increases base rates and causes the virtual machines to become more expensive.<br>• Base rates and virtual machine costs do not change with the utilization of the cluster.<br>• The difference between Usable Capacity after HA and Buffer and actual utilization is used to compute unallocated costs. |
| Cluster Actual Utilization | To calculate the base rates using the month to date average utilization of the cluster resources, select this option.<br>Base rates are calculated based on the total cost of the cluster and average utilization. Virtual machine costs are calculated from these base rates. Things to note:<br>• Lower utilization level causes base rates to be high and virtual machines also become more expensive.<br>• Higher utilization level causes base rates to be lower and virtual machines to become cheaper.<br>• Base rates and virtual machine costs can change frequently based on the utilization of the cluster.<br>• Unallocated cost of the cluster is near to zero.<br>• The costs for unused resources are distributed across all virtual machines based on their actual utilization within the cluster. |

4. Click **SAVE**.

## Publish Daily Cost Metrics for Virtual Machines

In VMware Aria OperationsVMware Cloud Foundation Operations, you can now publish daily cost metrics for all virtual machines. The daily cost metric of a virtual machine is the sum of daily cost of CPU, memory, storage, and additional cost associated with the virtual machine. Daily cost metrics provide granular details of the costs associated with the virtual machine.

**Formula to Calculate the Daily Cost and Monthly Cost of Virtual Machines**

You can calculate the daily cost associated with a virtual machine using the following formula.

| Virtual Machine Cost Elements | Calculation |
|---|---|
| Daily Total Cost of Virtual Machine | Daily total cost of virtual machine = Sum of Daily cost of (CPU + memory + storage + additional cost) |

The change in daily cost metrics also changes the way you calculate the effective month to date cost of a virtual machine. You can use the following formula to calculate Effective Month to Date cost for a virtual machine.

| Virtual Machine Cost Elements for a Month | Calculation |
|---|---|
| Effective MTD Cost of VM | Sum of CPU daily cost from the beginning of the month until now + Sum of memory daily cost from the beginning of the month until now + Sum of storage daily cost from the beginning of the month until now + Sum of additional daily cost from the beginning of the month until now |

**How to View the Daily Cost Metrics of a Virtual Machine**

To view the daily cost metrics of a virtual machine, from the menu, select **Inventory › vCenter Adapter**, select the specific **Virtual Machine**, and click the **Metrics** tab.

## Publish Tag Based Cost as Individual Metrics

You can publish tag-based additional cost as individual metrics using VMware Aria OperationsVMware Cloud Foundation Operations.To publish tag-based additional costs as individual metrics, you must first activate the Tag based Costing Metrics at the Global Settings level. If you activate the tag-based costing metrics at a VM level, then each of the tag-based cost is considered as an independent instance metric on the virtual machine.

**How to Activate Tag Based Costing Metrics**

From the left menu, click **Administration › Global Settings**, and then click **Cost/Price** and navigate to the **Tag based Costing Metrics**. Move the toggle button to the right to activate the Tag based costing metrics. Add a VM tag or custom properties based entries in Additional Cost Drivers. To know how to set the Additional Cost at a VM level, see Editing Additional Costs.

After you assign the cost for VMs with tags or custom properties, you must run the Cost Calculation and verify whether the additional cost is reflected in the costing metrics. To verify whether the tag-based cost metrics update is reflected in your cost, run the cost calculation and check the tag-based cost metrics. The additional cost metrics can be viewed from the following location, only for virtual machines which have the tags mentioned in the Additional Cost Driver: **Inventory › VM › Cost › Daily Tags and Custom Properties Cost**.

## Pricing Overview

You can create pricing cards in VMware Aria OperationsVMware Cloud Foundation Operations to calculate the price associated with your virtual infrastructure. You can assign pricing cards to vCenters or Clusters, depending on the pricing strategy determined by VMware Aria OperationsVMware Cloud Foundation Operations administrator. The pricing cards help you to set the price for each resource present in your virtual environment.

You can customize the pricing card as per your requirement. VMware Aria OperationsVMware Cloud Foundation Operations has two types of pricing cards, rate-based pricing card and cost-based pricing card. After configuring a pricing card, you can assign it to one or more vCenters or Clusters as determined by the pricing strategy.

> **NOTE**
> From this release, you can continue to create pricing cards or you can migrate vCenter pricing cards to vCenter pricing policies. For details about creating a new pricing card, see Add New Pricing Card. For details about migrating vCenter Pricing cards to vCenter pricing policies, see Migration of vCenter Pricing Cards to vCenter Pricing Policies.

**How Is Price Calculated**

In rate-based pricing policy VMware Aria OperationsVMware Cloud Foundation Operations calculates the virtual infrastructure price based on the rate card defined by you. For rate-based pricing policy VMware Aria OperationsVMware Cloud Foundation Operationslets you define cost elements as per your requirements.

The server recalculates the price every 24-hours, the price calculation for the new pricing cards is done in the next VMware Aria OperationsVMware Cloud Foundation Operations price calculation cycle.

### Hierarchy of Pricing Policy

The assignment of policy in VMware Aria OperationsVMware Cloud Foundation Operationswill be for Clusters and vCenters. The price is calculated for virtual machines, then it is aggregated and rolled up to vCenter. If there are two policies, a default policy for vCenter and another policy for Cluster, then the price calculation is based on the cluster policy for all the resources under the cluster. After that the cluster cost is rolled up to vCenter.

When a virtual machine is under VMware Aria Automation hierarchy and vCenter hierarchy, then the pricing is calculated based on the VMware Aria Automation hierarchy and the virtual machine is removed from the vCenter resources and included under VMware Aria Automation resources.

### Pricing Support for VMware Cloud on AWS Resources

You can create a pricing policy in VMware Aria OperationsVMware Cloud Foundation Operationsand assign it to VMware Cloud on AWS (VMC) resources, however you can only use the rate-based pricing policy for VMC-related objects.

> **NOTE**
> When you assign Cost-Based Policy for VMC resources, the policy is not applied, and price calculated for the policy is reported as zero.

### Add New Pricing Card

You can add and assign new pricing card to vCenter and Clusters in VMware Aria OperationsVMware Cloud Foundation Operations. The pricing card can be cost-based or rate-based, you can customize the cost-based pricing card and rate-based pricing card as per your requirement. After configuring the pricing card, you can assign it to one more vCenter or Clusters based on your pricing strategy.

1. From the left menu, click **Operations › Configuration › Policy Definitions**, and then select **Edit Policy**.
2. Click **VC Pricing** and then click **New Pricing Card** and configure the details of the pricing card.

**Table 146: Pricing Card Configuration**

| Parameter | Description |
|---|---|
| Name and Description | 1. Enter a name and description for your pricing card.<br>2. Optional: Select Default for Unassigned Workloads.<br>3. Click **Next**.<br>Default pricing card applies to all vCenter resources which do not have a direct cost policy assigned to them. |
| Basic Charges | Pricing is **Based on Cost** or **Based on Rate**. Select the type of pricing card.<br>For **Cost-based** pricing:<br>1. Enter the Cost Factor for the following.<br>a. CPU Cost<br>b. Memory Cost<br>c. Storage Cost<br>d. Additional Cost |

*Table continued on next page*

*Continued from previous page*

| Parameter | Description |
|---|---|
| | 2. Select the charging period as per your requirement, the options are **Hourly**, **Daily**, **Weekly**, and **Monthly**.<br>3. Select how to charge for the resources, the options are **Always** or **Only When Powered On**.<br>4. Click **Next**.<br>The cost is defined in VMware Aria Operations. If selected, a multiplication factor is required. For example, if you select 1.1 as a factor, the cost is multiplied by 1.1 resulting in a 10% increase to the calculated cost. The price equation using cost is: <cost> x <multiplication factor> = Price<br><br>**NOTE**<br>For accurate private cloud VM prices, you must edit the cost drivers, see Editing Cost Drivers and for accurate hybrid cloud provider VM prices, you must edit the cloud provider rate card, see Add or Edit Cloud Provider.<br><br>For **Rate-based** pricing:<br>1. Enter the CPU Rate in MHz per vCPU.<br>2. Enter the Memory Rate per GB.<br>3. Enter Storage Rate per GB.<br>4. Select the ChargingPeriod for all the values.<br>5. Select the Charge On Power State for all the values. |
| Guest OSes | 1. Enter the Guest OS Name.<br>2. Enter the base rate.<br>3. Select the charging period as per your requirement, the options are **Hourly**, **Daily**, **Weekly**, and **Monthly**. |
| Tags | Enter the Tag name and Tag Value. Define the charging method and base rate.<br><br>• Recurring - enter a base rate and define recurring interval as the charge period. The absolute rate value is required and it is added to the overall price.<br>• One time - define the one-time base rate charge. The absolute value is required and it is added as a one time price.<br>• Rate Factor - A multiplication factor is required that is applied to the select charge category.<br><br>Select how to charge the Tag based on powered on state. |

*Table continued on next page*

*Continued from previous page*

| Parameter | Description |
|---|---|
| Overall Charges | You can define overall charges to VMs that match this policy.<br>1. Enter the VM setup charges.<br>2. Enter the Recurring charge and select the time period from the drop-down menu. |
| Assignments | You can assign the new pricing card to vCenters and Clusters.<br>1. Select the vCenter or Cluster to which you want to apply the pricing card.<br>2. Click **Add** and Click **Finish**. |

The new pricing card details are displayed in the Pricing tab.

#### Migration of vCenter Pricing Cards to vCenter Pricing Policies

When you upgrade to this version and later versions of VMware Aria OperationsVMware Cloud Foundation Operations, and you have vCenter Pricing Cards, you can migrate them to VMware Aria OperationsVMware Cloud Foundation Operations policies to utilize the full capability of VMware Aria OperationsVMware Cloud Foundation Operations policies.

**Steps to Migrate vCenter Pricing Cards to vCenter Pricing Policies**

1. From the left menu, click **Operations › Configuration**, click **Cost Drivers**, and then click the **Pricing** tab.
2. From the text box, click the **Actions** drop-down and then click **Migrate Now**.

After pricing card migration is complete, the following scenarios are available:

- **Pricing Policies Migration:** All the pricing rate cards will be converted to separate policies under the default policy. The converted policies will have the same name as the rate cards with the rate card UUID added as a suffix.
- **Pricing Cards Assignment Migration:**
  - If there is just one active policy before migration, then all the rate card assignments will be transferred to the newly created policies.
  - If there is more than one active policy you will have to manually reassign policies to the relevant objects after the migration. It is recommended to proceed with policy assignments immediately after the migration to ensure continuous pricing calculations.
    The rate cards will no longer be available. The new updated capability will be accessible under each policy under Policy Definition in VMware Aria OperationsVMware Cloud Foundation Operations.
  - Default rate card configuration will be migrated to the policy. However, it will not be set as the default. You can achieve this by either making it the default policy or by utilizing VMware Aria OperationsVMware Cloud Foundation Operations Custom Groups capabilities.

After migration of the vCenter pricing cards, you can view the **VC Pricing Migration Audit** logs from **Migration › Audit**, and then click the **VC Pricing Migration Audit** tab. For more details, see VC Pricing Migration Audit.

## Cost Calculation Status Overview

You can check the ongoing status of manually triggered cost calculation process.

Cost calculation by default, occurs daily and whenever there is a change in the inventory or cost drivers values. You can trigger the cost calculation manually so that changes in the inventory and cost driver values reflect accordingly on the VM

cost without having to wait there for any failures in the cost calculation process. It also shows default schedules time for next cost calculation process.

> **NOTE**
> To run the manual cost calculation, from the left menu click **Administration › Control Panel**, and then click **Cost Calculation › Run**.

#### Migration of Cost Driver Configuration from vRealize Business for Cloud to VMware Aria OperationsVMware Cloud Foundation Operations

vRealize Business for Cloud supports migration of cost driver configuration from vRealize Business for Cloud to VMware Aria OperationsVMware Cloud Foundation Operations. You can migrate cost driver configuration from vRealize Business for Cloud 7.x or later to vRealize Operations 6.7 or 7.5.

For more information about the migration process, see the KB article https://kb.vmware.com/s/article/55785.

## Costing Enhancements

In VMware Aria OperationsVMware Cloud Foundation Operations, a new global property Cluster Utilization Ceiling Factor is introduced. Using Cluster Utilization Ceiling Factor, you can specify the ceiling value and calculate the base rate for a cluster.

You can use the ceiling factor only if the base rate cost calculation is done using Cluster Actual Utilization method. After you set the ceiling factor value, the Actual Utilization of the cluster is rounded off to the next available multiple of the ceiling value. When ceiling value is 0, Expected Utilization is equal to actual utilization. When ceiling value is 20, it is not considered as special case, actual utilization is rounded off to the next multiple.

> **NOTE**
> The ceiling value range is from 0 to 20. If the number is out of this range, the default value of five is used as the ceiling number.

#### How to Set the Cluster Base Rate Calculation Method

To change the Cluster Base Rate Calculation method, you must go to **Operations › Configuration**, and then click **Cost Drivers › Cluster Cost**. Click **Change** next to the Cluster Base Rate calculation method and select Cluster Actual Utilization.

#### Where to Find Cluster Utilization Ceiling Factor

To set the ceiling value for a cluster, you must go to **Administration › Global Settings**, and then click **Cost/Price › Cluster Utilization Ceiling Factor**. Enter the ceiling value between 0 and 20 and click **Save**.

To view the change is cost metrics, run the Cost Calculation Status and select a cluster .
If the Actual Utilization of the cluster for CPU is 30 % and Memory is 45%, and the ceiling value specified is 10, then

- Cluster Expected CPU Utilization (%) = 40
- Cluster Memory Expected Utilization (%) = 50

Actual Cluster Utilization is rounded off to the ceiling value.
If you set the Cluster Utilization Ceiling Factor to either 0 or 20, then the value of Expected Memory Utilization changes to the next number. For example, if you set the ceiling factor to 0 then, the expected utilization value changes to 1.

#### Support to Roll up Name Space Cost Metrics

The cost metrics of Point of Delivery (Pod) virtual machines (VMs) has been enhanced to support the following scenarios:

- Cost metrics of Pod VMs are rolled up to the Name Space and Guest Cluster level.
- All the cost metrics of VMs, Pods, and guest cluster which are present under Name Space are rolled up to Name Space and Guest Cluster level.

| Old Cost Metrics | Rolled up Cost Metrics |
|---|---|
| Effective MTD Total Cost | Aggregate Additional Daily Cost |
| Deleted VM Daily Cost | Aggregate Deleted VM Daily Cost |
| Daily CPU Cost | Aggregate CPU Daily Cost |
| Daily Memory Cost | Aggregate Memory Daily Cost |
| Daily Storage Cost | Aggregate Storage Daily Cost |
| Daily Additional Cost | Aggregate Additional Daily Cost |

**Reclaimable Hosts Cost Metric**

You can use the cost metrics at cluster level to identify the clusters with reclaimable hosts and the potential cost savings from reclaiming these hosts. To know the cost associated with all the reclaimable hosts in a cluster, check the value of Total Host Reclaimable Host Cost metric.

**How to View Reclaimable Host Cost**

To view the reclaimable host cost, go to **Inventory › vCenter Adapter**, and then click **Cluster Compute Resource › Cost**.

You can also view the total host reclaimable cost using **Inventory › vCenter Adapter**, click **vCenter Adapter › vSphere World**, and then click **Metrics › Cost**.

> **NOTE**
> If the cluster does not have reclaimable hosts, then the cost metric associated with the reclaimable host is not displayed.

**Realized Cost Savings Using Reclamation Suggestion**

In VMware Aria OperationsVMware Cloud Foundation Operations you can track the cost savings using reclamation suggestions. Using the reclamation option, you can view the cost, capacity, and allocation metrics related to individual data centers. The metrics provides an estimate of the potential savings achieved through VMware Aria OperationsVMware Cloud Foundation Operations.

You can track the realized cost savings and actual capacity reclaimed for data centers, in the following scenarios.

- Reclaim the cost for Idle VMs by deleting the VM.
- Reclaim the cost for Powered off VMs by deleting the VM.
- Reclaim the cost for Idle VMs by powering off the VM.
- Reclaim the cost for snapshots VMs by deleting the snapshot.
- Reclaim the cost for orphaned disks by deleting the orphaned disk space.
- Reclaim the cost by removing vCPU and Memory from an oversized VM.
- Reclaim the cost by removing a host from the vCenter.

**Costing for Oversized VM and Undersized VM**

Rightsizing is defined as changing the amount of resources allocated to a VM based on the Recommended Size for a VM. Recommended Size is the maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The warning threshold is the period during which the time remaining is green. If the warning threshold value for time remaining is 120 days, which is the default value, the recommended size is the maximum projected utilization 150 days into the future. While rightsizing a VM can reclaim capacity, the change in allocation may not equal the amount of reclaimable capacity.

**Quantifying the Effect on Capacity Due to Rightsizing**

**Demand Model**

- **Reclaimable CPU Usage (GHz)**: If an oversized VM's CPU usage is 100MHz before rightsizing, removing vCPU's will not change its CPU usage and it should still be at 100MHz. This means there is no reclaimable capacity associated with overallocation of vCPUs. Reclaimable CPU Usage for oversized VM's will always be 0 MHz.
- **Reclaimable Memory Consumed (GB)**: An oversized VM can have reclaimable memory only if consumed memory is greater than the new recommended size of the VM. The reclaimable memory capacity is the difference between consumed memory and recommended size.
- **Increased CPU Usage (GHz)**: CPU usage of an undersized VM is expected to be the current CPU Demand. The difference between CPU Demand and CPU Usage is the expected increase in capacity utilized after rightsizing.
- **Increased Memory Consumed (GB)**: It can be expected for consumed memory to increase by the same amount of memory recommended to add to an undersized VM.

**Allocation Model**

In case of allocation model, you can directly pick the recommendation provided which is given as a part of the metric groups **Summary|Oversized** and **Summary|Undersized**.

**Potential Cost Savings Calculation Detail**

- **Oversized CPU Utilization**: $0 since Reclaimable CPU Usage (GHz) is always 0.
- **Oversized Memory Utilization**: Reclaimable Memory Consumed (GB) * Cluster Memory Base Rate.
- **Oversized CPU Allocation**: vCPU(s) to Remove * Allocation Cluster CPU Base Rate.
- **Oversized Memory Allocation**: Memory to Remove * Allocation Cluster Memory Base Rate.

**Potential Cost Increase Calculation Detail**

- **Undersized CPU Utilization**: Increased CPU Usage (GHz) * Cluster CPU Base Rate.
- **Undersized Memory Utilization**: Increased Memory Consumed (GB) * Cluster Memory Base Rate.
- **Undersized CPU Allocation**: vCPU(s) to Add * Allocation Cluster CPU Base Rate.
- **Undersized Memory Allocation**: Memory to Add * Allocation Cluster Memory Base Rate.

The rightsizing value calculated here is available as part of
- **Potential Savings** metric (For VM) for Oversized VMs.
- **Potential Increase** metric (For VM) for Undersized VMs.

> **NOTE**
> Reclaimable memory consumed, Increased CPU Usage, and Increased memory consumed are metrics that are available for reference under **Summary|Oversized** metrics and **Summary|Undersized** metrics respectively.

## Viewing and Configuring Compliance

VMware Aria Operations helps you visualize the compliance of the objects in your inventory by displaying compliance benchmarks which are measured against a set of standard rules, regulatory best practices, or custom alert definitions. You can discover non-compliant objects and the associated compliance violations and take remedial measures.

## Measuring Compliance of Objects

VMware Aria Operations continuously monitors your infrastructure to ensure that it remains in compliance. Compliance is an ongoing process. VMware Aria Operations evaluates the collected data against the defined policies. It assigns a compliance score to each object or group based on how well they adhere to the policies. The compliance score is typically represented as a percentage.

## How Compliance Benchmarks Work

Compliance is used to monitor the vCenter Server instances, NSX, vSAN and Cloud Environments. Objects which can be monitored include the hosts, virtual machines, distributed port groups, and distributed switches in your environment. Compliance benchmarks helps to ensure that the settings on your objects meet the defined standards. Score cards help you proactively detect compliance problems in objects manged by VMware Aria Operations. The compliance benchmarks are measured against a set of standard rules, regulatory best practices, or custom alert definitions.

All the compliance standards in VMware Aria Operations, including any standards that you define, are based on alert definitions. Only alert definitions of the Compliance subtype are counted. Custom score cards can monitor user-defined alerts. The alerts and symptom definitions are based on the properties and metrics of the underlying object.

When VMware Aria Operations detects non-compliance with a policy, it can generate alerts or notifications. Depending on the severity of the non-compliance, you can configure automated remediation actions to bring the object back into compliance.

You can manage all compliance related tasks from the **Operations › Compliance** page. The data sources are displayed in a carousel on the top of the page. To see a compliance score card, you must first configure the data source that VMware Aria Operations can monitor, and then activate the benchmarks for those types of data sources. When you activate a benchmark for a data source, you select an applicable policy. VMware Aria Operations then activates the appropriate alert definitions in the policy to measure compliance.

## Data Sources for Calculating Compliance

VMware Aria Operations can measure compliance from different data sources to ensure that your virtualized environment adheres to predefined policies and standards. These data sources provide the information necessary to evaluate the compliance of your infrastructure. The data sources that VMware Aria Operations can use for compliance measurement are as follows:

- vCenter Systems
- VMware Cloud Foundation
- Google Cloud VMware Engine

**NOTE**
For VMware Aria Operations to measure compliance against these data sources, you must first configure them. See the relevant topic in the Integrating Data Sources with VMware Aria OperationsVMware Cloud Foundation Operations chapter of the *Configuring VMware Aria Operations* guide.

**NOTE**
In version 8.16, the symptom set for the FISMA Security Standards, DISA Security Standards and vSphere compliance pack were updated to generate an alert when the following conditions are false:

- Block Override Allowed should be true
- Port Config Reset at Disconnect should be true

**NOTE**
In version 8.16, a bug involving DVPG symptoms on Promiscuous mode, MAC address changes, Forged transmits not considering the uplink/non-uplink state was fixed.
The symptoms on DVPG properties **`allow_promiscuous`**, **`forged_transmits`** and **`mac_changes`** are defined in the following compliance packs:

- CIS Security Standards
- DISA Security Standards
- FISMA Security Standards
- HIPAA
- ISO Security Standards
- PCI DSS Compliance Standards
- vSphere Security Configuration Guide

Compliance benchmarks on VMware Cloud on AWS, VMware Cloud Foundation, Oracle Cloud VMware Solution, Azure VMware Solution, and Google Cloud VMware Engine are applicable only on customer VMs that you have deployed in the respective data centers. For more details on these integrations, see the VMware Aria Operations for Integrations Product Documentation, or Integrating Data Sources with VMware Aria OperationsVMware Cloud Foundation Operations.

You can automate the remediation of some of the alerts by installing the Management Pack for VMware Aria Automation Orchestrator. See the management pack documentation in the *VMware Aria Operations for Integrations Product Documentation* for more details.

## Compliance Benchmarks

### VMware SDDC and Benchmarks

Displays score cards based on alerts which are measured against the latest hardening guides:

- vSphere Security Configuration Guide
- vSAN Security Configuration Guide
- NSX Security Configuration Guide

For more details, see VMware SDDC Benchmark Details.

### VMware Cloud Foundation Benchmarks

Displays score cards based on alerts which are measured in VMware Cloud Foundation domains based on the following:

- VCF 4.2 Audit Guide
- VCF 4.3 Audit Guide
- VCF 4.4 Audit Guide
- VCF 4.5 Audit Guide
- VMware Cloud Foundation Operations CSA Compliance Pack for VMware Cloud Foundation. For details of the conditions implemented in VMware Aria Operations, see the Knowledge Base article 371288.

  **NOTE**
  You must install and activate the VMware Aria Operations CSA Compliance Pack for VMware Cloud Foundation after downloading the .PAK file from Marketplace

The alerts are based on objects in your VMware Cloud Foundation environment.

For more details, see the topic, VMware Cloud Foundation Benchmarks based on VMware Cloud Foundation Compliance Kits.

### Custom Benchmarks

Displays benchmarks that you define. Use compliance alerts from integrations, and regulatory management packs, or define your own alerts to monitor. You can define up to five custom score cards. You can import custom benchmarks from other instances of VMware Aria Operations.

### Regulatory Benchmarks

Displays benchmarks for industry standard regulatory compliance requirements. You can install compliance packs for the following regulatory standards:

- Health Insurance Portability and Accountability Act (HIPAA)
- Payment Card Industry Data Security Standard (PCI DSS) Compliance Standards
- CIS Security Standards
- Defense Information Systems Agency (DISA) Security Standards
- The Federal Information Security Management Act (FISMA) Security Standards
- International Organization for Standardization (ISO) Security Standards

For more details, see Regulatory Benchmark Details.

For instructions on installing these compliance packs, see Activate a Regulatory Benchmark.

**VMware Cloud Foundation Benchmarks based on VMware Cloud Foundation Compliance Kits**

Security and Compliance for VMware Cloud Foundation provides guidance for auditors who would want to evaluate their VMware Cloud Foundation environment. You can leverage SDDC benchmarks and regulatory frameworks available under VMware Cloud Foundation to be able to assess individual SDDC components based on needed benchmarks or regulatory frameworks of your choice.

**VCF Audit Guide**

Compliance Kit for VMware Cloud Foundation is a solution that builds on top of VMware Cloud Foundation and leverages security fundamentals. The kit addresses the top ten most frequently requested compliance standards, regulations, and frameworks.

The compliance kit is designed and validated to tailor security configurations without impacting the ability of VMware Cloud Foundation to meet its design objectives. The kit can assist organizations to secure information systems in a compliance context.

The VCF Audit Guide is part of the Compliance Kit for VMware Cloud Foundation and can be used to evaluate both default and non-default configurations. For more information, see the topic, Compliance Kit for VMware Cloud Foundation in the *VMware Cloud Foundation Product Documentation*.

The .xlsx file, that is based on the VCF Audit Guide, containing a collection of all compliance conditions required for assessing compliance of VCF environments is available in the KB article: KB94849.

Versions of the supported VMware Cloud Foundation Compliance kits are:
- VCF 4.2 Audit Guide
- VCF 4.3 Audit Guide
- VCF 4.4 Audit Guide
- VCF 4.5 Audit Guide

VMware Aria Operations supports the following products for VMware Cloud Foundation Benchmarks:
- ESXi,
- SDDC Manager,
- vCenter Server,
- vSAN
- NSX

Within these products, VMware Aria Operations can measure the compliance of the following resources:
- VMware Distributed Virtual Switch
- VMware Adapter Instance
- Virtual Machine
- Host System
- Logical Router
- NSX
- NSX Manager Service
- Logical Switch
- Management Cluster
- VCF Adapter Instance
- vSAN
- DC
- Cluster

**NOTE**
Assessment of objects starting from version VMware Cloud Foundation 4.5 and above is done based on VCF 4.5 Audit Guide.

**NOTE**
You must configure the VMware Cloud Foundation data source before you can configure the VCF benchmarks. For information, see the topic, VMware Cloud Foundation in the *Configuring VMware Aria Operations* guide.

## VMware Aria Operations CSA Compliance Pack for VMware Cloud Foundation

The Cloud Security Alliance (CSA) framework holds significant importance within the sphere of cloud computing, offering comprehensive guidance and best practices for ensuring the security of cloud-based systems and services.

The VMware Cloud Foundation Operations CSA Compliance Pack for VMware Cloud Foundation is available as a Management Pack. Before you can enable the VMware Aria Operations CSA Compliance Pack for VMware Cloud Foundation, you must install and activate the integration. For more details, see the VMware Aria Operations for Integrations Documentation.

### VMware SDDC Benchmark Details

VMware SDDC (Software-Defined Data Center) security benchmarks are guidelines or standards that provide recommendations and best practices for securing VMware's Software-Defined Data Center infrastructure. These benchmarks help organizations assess and enhance the security of their virtualized data center environments.

## vSphere Security Configuration Guide

For details of the conditions that allow automated compliance assessment and the list of controls that can be used to perform manual checks based on the VMware vSphere 7 Security Configuration Guide Update 3 and VMware vSphere 8 Security Configuration Guide, see KB 94171.

For details of the conditions that allow automated compliance assessment and the list of controls that can be used to perform manual checks based on the VMware vSphere 7 Security Configuration Guide Update 3 and VMware vSphere 8 Security Configuration Guide, see KB 88721.

The vSphere Security Configuration Guide is a comprehensive document provided by VMware that outlines best practices and recommendations for securing VMware vSphere, which is a virtualization platform used for creating and managing virtualized data centers.

The guide covers various aspects of vSphere security and provides configuration guidelines to help administrators protect their vSphere infrastructure from potential threats and vulnerabilities. It addresses security considerations for different components of the vSphere environment, including ESXi hosts, vCenter Server, virtual machines, networking, and storage.

The vSphere Security Configuration Guide typically covers the following areas:

- ESXi Host Hardening: It provides recommendations for securing the ESXi hosts, including configuring authentication and authorization settings, disabling unnecessary services, enabling security features like Secure Boot and Trusted Platform Module (TPM), and implementing network security measures.
- vCenter Server Hardening: The guide offers best practices for securing the vCenter Server, which is a centralized management platform for vSphere. It covers securing the vCenter Server installation, securing access and authentication, configuring roles and permissions, and enabling auditing and logging.
- Virtual Machine Security: It provides recommendations for securing virtual machines (VMs), such as implementing secure VM configurations, using virtual machine encryption, enabling guest operating system security features, and protecting VMs from unauthorized access.

- Networking Security: The guide offers guidelines for securing the vSphere networking infrastructure, including configuring virtual switches, VLANs, and firewalls, implementing network segmentation, and securing network communication between vSphere components.
- Storage Security: It covers best practices for securing the vSphere storage environment, including securing access to storage devices, implementing storage encryption, and protecting virtual machine data.
- Monitoring and Auditing: The guide emphasizes the importance of monitoring and auditing vSphere components, and provides recommendations for enabling and configuring logging, using security information and event management (SIEM) tools, and monitoring for suspicious activities.

The vSphere Security Configuration Guide is regularly updated by VMware to incorporate the latest security recommendations and considerations. It serves as a valuable resource for administrators seeking to secure their vSphere infrastructure and protect their virtualized environments from potential security risks.

**vSAN Security Configuration Guide**

The vSAN Security Configuration Guide is a document provided by VMware that offers guidance and best practices for securing VMware vSAN (Virtual SAN), which is a software-defined storage solution integrated into the vSphere environment.

The guide focuses on helping administrators secure the vSAN infrastructure and protect the data stored within it. It provides recommendations for various security considerations related to vSAN, including authentication, encryption, network security, and access controls.

Here are some key areas covered in the vSAN Security Configuration Guide:

- Authentication and Authorization: The guide provides recommendations for configuring strong authentication mechanisms for vSAN components, such as vCenter Server and ESXi hosts. It suggests using secure protocols, enforcing secure password policies, and implementing multi-factor authentication where possible. It also covers role-based access control (RBAC) and permissions management.
- Encryption: vSAN supports encryption at rest, which ensures that data stored on vSAN disks is protected. The guide offers guidance on configuring vSAN encryption, including selecting the appropriate encryption algorithm, managing key management servers (KMS), and enabling encryption for data at rest.
- Network Security: It provides recommendations for securing network communication within the vSAN environment. This includes using secure protocols, enabling network encryption (VMkernel Encryption), and configuring firewall rules to control network traffic between vSAN components.
- Auditing and Logging: The guide emphasizes the importance of monitoring and auditing vSAN operations. It offers recommendations for enabling and configuring logging, implementing log management solutions, and monitoring vSAN events for potential security incidents.
- Secure Configuration Settings: It provides a list of vSAN-specific configuration settings that can enhance security. This includes recommendations for enabling features like Data-at-Rest Encryption, Secure Erasure, and ensuring proper network configurations.
- Compliance and Hardening: The guide addresses compliance considerations and provides information on how vSAN aligns with various security frameworks, such as the Center for Internet Security (CIS) benchmarks. It also offers recommendations for hardening vSAN components, including ESXi hosts and vCenter Server.

The vSAN Security Configuration Guide is regularly updated by VMware to incorporate the latest security best practices and recommendations. It serves as a valuable resource for administrators who want to ensure the security of their vSAN deployments and safeguard their storage infrastructure and data.

**NSX Security Configuration Guide**

The VMware Cloud Foundation Operations Compliance Pack for NSX-T is updated to support the following standards:

- NSX-T 3.2 Security Configuration Guide
- NSX-T 3.1 Security Configuration Guide
- NSX-T 3.0 Security Configuration Guide

- **NOTE**
  NSX-T is renamed to NSX.

For more details, see KB 93136.

The NSX Security Configuration Guide is a comprehensive document provided by VMware that offers guidance and best practices for securing VMware NSX, which is the next-generation network and security platform for software-defined networking.

The NSX Security Configuration Guide focuses on helping administrators secure the NSX infrastructure and protect network traffic, ensuring the confidentiality, integrity, and availability of data. It provides recommendations for various security considerations related to NSX, including authentication, authorization, network security, encryption, and compliance.

Here are some key areas covered in the NSX Security Configuration Guide:
- NSX Component Hardening: The guide provides recommendations for securing NSX components such as NSX Manager, NSX Controllers, and NSX Edge nodes. It includes guidelines for configuring strong authentication, implementing role-based access control (RBAC), and securing administrative access to NSX components.
- Network Security: It covers best practices for implementing network security within NSX. This includes recommendations for creating and managing security groups, implementing distributed firewalling, configuring security policies, and enforcing microsegmentation to control network traffic.
- Authentication and Authorization: The guide offers recommendations for configuring strong authentication mechanisms for NSX, such as integrating with identity providers, enabling multi-factor authentication, and defining access controls based on user roles and permissions.
- Secure Network Communication: It provides guidance for securing network communication within the NSX environment. This includes configuring Transport Layer Security (TLS) encryption for communication between NSX components, securing communication between NSX and vCenter Server, and protecting east-west and north-south traffic.
- Virtual Machine Security: The guide provides recommendations for securing virtual machines (VMs) deployed in NSX environments. This includes implementing security groups, defining security policies, and leveraging NSX's integration with third-party security solutions for advanced VM protection.
- Logging and Auditing: It emphasizes the importance of logging and auditing in NSX environments. The guide offers recommendations for enabling and configuring logging, integrating NSX with logging and monitoring solutions, and monitoring for security events and anomalies.
- Compliance and Hardening: The guide addresses compliance considerations and provides information on how NSX aligns with various security frameworks and regulations. It also offers recommendations for hardening NSX components and ensuring compliance with security policies.

The NSX Security Configuration Guide is regularly updated by VMware to incorporate the latest security best practices and recommendations. It serves as a valuable resource for administrators who want to ensure the security of their NSX deployments, implement effective network security controls, and protect their software-defined networking infrastructure.

**Regulatory Benchmark Details**

Regulatory compliance benchmarks are standards or guidelines that help organizations measure and assess their level of compliance with applicable laws, regulations, and industry standards.

**Health Insurance Portability and Accountability Act (HIPAA)**

The Health Insurance Portability and Accountability Act (HIPAA) is a U.S. federal law enacted in 1996. HIPAA establishes standards and regulations to protect the privacy, security, and confidentiality of individuals' personal health information (PHI) and electronic health records (EHRs) in the healthcare industry.

The HIPAA Privacy Rule and Security Rule are two key components of the HIPAA standard:
- HIPAA Privacy Rule: The Privacy Rule sets standards for the use and disclosure of PHI by covered entities, which include healthcare providers, health plans, and healthcare clearinghouses. It grants individuals certain rights over

their health information, such as the right to access, request amendments, and obtain an accounting of disclosures. Covered entities are required to implement safeguards to protect PHI, provide patients with notice of privacy practices, and obtain written authorization for certain uses and disclosures of PHI.
- HIPAA Security Rule: The Security Rule establishes security standards for protecting electronic PHI (ePHI). It requires covered entities and their business associates to implement administrative, physical, and technical safeguards to ensure the confidentiality, integrity, and availability of ePHI. These safeguards include risk assessments, access controls, encryption, audit controls, disaster recovery plans, and employee training on security awareness.

The HIPAA standard applies to healthcare providers, health plans, healthcare clearinghouses, and their business associates who handle PHI or ePHI. Compliance with HIPAA regulations is mandatory, and non-compliance can result in significant penalties, including financial fines and potential criminal charges.

HIPAA also includes provisions related to the electronic exchange of health information and establishes the Health Information Technology for Economic and Clinical Health (HITECH) Act, which promotes the adoption and meaningful use of electronic health records.

### Payment Card Industry Data Security Standard (PCI DSS) Compliance Standards

The Payment Card Industry Data Security Standard (PCI DSS) is a set of compliance standards established by the major payment card brands, including Visa, Mastercard, American Express, Discover, and JCB. PCI DSS aims to ensure the security of cardholder data and protect against fraud and unauthorized access within the payment card industry.

The PCI DSS compliance standards consist of twelve high-level requirements, organized into six control objectives. These requirements outline security measures that organizations handling payment card data must implement:
- Build and Maintain a Secure Network
- Protect Cardholder Data
- Maintain a Vulnerability Management Program.
- Maintain an Information Security Policy

PCI DSS compliance requirements vary based on the organization's level of involvement in payment card transactions, classified as levels 1 to 4. Level 1 merchants and service providers with the highest transaction volumes have more stringent requirements and undergo annual on-site audits by a Qualified Security Assessor (QSA). Level 2, 3, and 4 merchants may have different validation requirements, ranging from self-assessment questionnaires (SAQs) to external vulnerability scans.

Compliance with PCI DSS is necessary for entities involved in payment card processing, including merchants, service providers, and payment processors. Non-compliance may result in penalties, fines, increased transaction fees, or restrictions on card acceptance.

It's important to note that while this information provides an overview of the PCI DSS compliance standards, specific requirements and guidance can evolve over time. Therefore, it is recommended to consult the official PCI Security Standards Council (PCI SSC) website and the latest PCI DSS documentation for the most up-to-date information and requirements.

### CIS (Center for Internet Security) Security Standards

The VMware Aria Operations Compliance Pack for CIS is updated to support the following benchmarks:
- CIS_VMware_ESXi_6.7_Benchmark_V1.3.0
- CIS_VMware_ESXi_7.0_Benchmark_V1.2.0

For more details, see the Knowledge Base article: 371283.

The CIS (Center for Internet Security) Security Standards are a set of best practices and guidelines for securing computer systems and networks. The CIS organization is a non-profit entity that collaborates with experts from various industries to develop and promote consensus-based security configurations and benchmarks.

The CIS Security Standards include two primary components:

- CIS Controls: The CIS Controls are a set of 20 security actions that organizations can take to mitigate the most common and impactful cyber threats. These controls are prioritized based on their effectiveness in reducing risk. They cover various security domains, including asset management, access control, incident response, network security, and security awareness training. The CIS Controls are regularly updated to address emerging threats and evolving technology landscapes.
- CIS Benchmarks: CIS Benchmarks provide detailed configuration guidelines for securing specific technology platforms and systems. These benchmarks outline recommended settings and configurations for operating systems, applications, and network devices to ensure security and reduce vulnerabilities. CIS Benchmarks are created through a consensus-driven process involving input from cybersecurity experts, vendors, and practitioners.

CIS Security Standards are known for their practical and actionable nature, providing step-by-step instructions and specific configuration recommendations. They are widely adopted across industries and are used as a reference by organizations to assess, improve, and maintain the security of their IT systems and networks.

The CIS organization regularly updates its security standards and benchmarks to address emerging threats, technology advancements, and changes in regulatory requirements. The CIS Security Standards are available to the public, and organizations can leverage them as a valuable resource for enhancing their cybersecurity posture and reducing the risk of cyberattacks.

**Defense Information Systems Agency (DISA) Security Standards**

The Defense Information Systems Agency (DISA) establishes and provides security standards and guidelines for the U.S. Department of Defense (DoD) and its information systems. DISA is responsible for ensuring the secure operation and defense of DoD's global information infrastructure.

The VMware Aria Operations Compliance Pack for Defense Information Systems Agency (DISA) Security Standards is updated to the following version in VMware Aria Operations 8.18:

- VMware vSphere 7.0 Version 1 Release 2: 26 July 2023
- VMware vSphere 8 Version 1 Release 1: 03 November 2023

For more details, see the Knowledge Base article: 371287

DISA has developed several security standards and guidelines to protect sensitive information and ensure the integrity, availability, and confidentiality of DoD systems. Some of the key security standards and guidelines provided by DISA include:

- Security Technical Implementation Guides (STIGs): STIGs are a set of guidelines and configuration standards for various operating systems, applications, and network devices. They provide detailed instructions on how to secure and configure these systems to meet DoD security requirements. STIGs cover a wide range of technologies, including Windows, Linux, Cisco devices, databases, and web servers.
- Security Requirements Guides (SRGs): SRGs are comprehensive documents that outline security requirements for specific technology platforms, systems, or applications. They provide guidance on how to secure and configure systems in accordance with DoD security policies. SRGs address various security domains, including access control, identification and authentication, audit and accountability, and encryption.
- Security Technical Implementation Guides (STIGs) Viewer: DISA provides a STIG Viewer tool that helps organizations assess and implement STIG recommendations. The STIG Viewer automates the process of checking system configurations against STIG requirements, allowing organizations to identify and remediate security vulnerabilities more efficiently.
- Information Assurance Vulnerability Management (IAVM): DISA maintains the IAVM program, which identifies and manages vulnerabilities in DoD systems. IAVM alerts provide timely information about security vulnerabilities and patches. Organizations within the DoD are required to promptly apply these patches to mitigate potential risks.
- DoD Cybersecurity Discipline Implementation Plan (CDIP): The CDIP outlines the implementation and management of cybersecurity practices within the DoD. It provides guidelines and best practices for managing risks, protecting systems, responding to incidents, and fostering a culture of cybersecurity awareness.

DISA's security standards and guidelines play a critical role in ensuring the security and resilience of DoD systems and information assets. They are constantly updated and refined to address emerging threats and align with evolving

cybersecurity practices. Organizations within the DoD are expected to adhere to these standards to maintain the security of their systems and networks.

**The Federal Information Security Management Act (FISMA) Security Standards**

The Federal Information Security Management Act (FISMA) is a U.S. federal law enacted in 2002. FISMA establishes a framework for securing information systems and managing cybersecurity risks within federal government agencies and their contractors. FISMA requires federal agencies to develop, implement, and maintain information security programs to protect sensitive government information.

While FISMA itself does not provide detailed security standards, it sets requirements for federal agencies to follow certain security guidelines and standards, including those established by the National Institute of Standards and Technology (NIST). NIST Special Publication (SP) 800-53, titled "Security and Privacy Controls for Federal Information Systems and Organizations," is a key document referenced under FISMA.

NIST SP 800-53 provides a catalog of security controls that federal agencies must implement to protect their information systems. The controls cover various areas, including access control, incident response, configuration management, encryption, network security, and security assessment and authorization. The controls are categorized into families and are tailored to address specific security requirements.

FISMA requires federal agencies to develop and maintain a risk-based approach to information security. This involves conducting risk assessments, implementing security controls based on the identified risks, periodically testing and evaluating the effectiveness of these controls, and ensuring continuous monitoring of information systems.

Under FISMA, federal agencies are also required to undergo annual security assessments, including independent audits, to evaluate the effectiveness of their information security programs and controls. The results of these assessments are reported to the Office of Management and Budget (OMB) and Congress.

FISMA compliance is crucial for federal agencies to demonstrate their commitment to protecting government information and ensuring the security of their information systems. It helps establish a standardized approach to information security across federal government entities and aligns with other security frameworks and standards, such as the NIST Cybersecurity Framework and NIST Risk Management Framework.

It's important to note that FISMA requirements may evolve over time, and agencies should refer to the latest guidance provided by NIST and other authoritative sources to ensure compliance with FISMA security standards.

**International Organization for Standardization (ISO) Security Standards**

The International Organization for Standardization (ISO) is an independent, non-governmental international standardization body that develops and publishes international standards across various industries. ISO has also created a series of security standards specifically related to information security management systems (ISMS). The most well-known among them is ISO/IEC 27001.

ISO/IEC 27001: The ISO/IEC 27001 standard specifies the requirements for establishing, implementing, maintaining, and continuously improving an ISMS within the context of an organization. It provides a systematic and risk-based approach to managing the security of sensitive information. The standard covers areas such as risk assessment, information security policies, asset management, access control, incident management, and compliance. ISO/IEC 27001 is widely adopted by organizations globally and serves as a benchmark for information security management.

ISO/IEC 27002: ISO/IEC 27002 (formerly known as ISO/IEC 17799) is a code of practice for information security controls. It offers guidance and recommendations for implementing security controls and safeguards based on the best practices of information security management. ISO/IEC 27002 covers a broad range of security areas, including organizational security, human resource security, physical and environmental security, communications and operations management, and compliance.

ISO/IEC 27005: ISO/IEC 27005 provides guidelines for conducting risk assessments in the context of information security. It offers a structured approach for identifying, analyzing, evaluating, and treating information security risks. ISO/IEC 27005

helps organizations assess the potential impact of risks, determine risk tolerance, and make informed decisions on implementing appropriate security controls.

ISO/IEC 27017 and ISO/IEC 27018: These standards specifically focus on cloud security. ISO/IEC 27017 provides guidelines for implementing information security controls in cloud computing environments, while ISO/IEC 27018 offers guidance for protecting personal data in the cloud and addresses privacy concerns related to cloud services.

ISO/IEC 27701: This standard is an extension to ISO/IEC 27001 and provides guidelines for implementing a Privacy Information Management System (PIMS). ISO/IEC 27701 helps organizations establish and maintain controls to protect personal data and comply with privacy regulations, such as the General Data Protection Regulation (GDPR).

ISO security standards provide a framework for organizations to establish effective information security management practices. Compliance with these standards demonstrates a commitment to securing sensitive information, managing risks, and implementing robust security controls. Organizations can seek ISO certification through a formal audit process conducted by accredited certification bodies to validate their adherence to ISO security standards.

### Compliance Score Cards

The compliance page in VMware Aria OperationsVMware Cloud Foundation Operations displays score cards for each type of benchmark. A score card is a compliance visualization term.

#### What is a Compliance Score Card

Score cards in the Compliance landing page display the number of non-compliant objects, and the total number of objects affected by each hardening guide and the compliance score which is counted as the ratio of compliant objects to total number of objects assessed by the given benchmark, represented in percentage. In addition, you can see the breakdown of the total number of objects that are compliant and non-compliant. You can click a score card to see more details, including alerts that were triggered based on the compliance standards.

The compliance score card of an object is counted as the smallest rounded off integer (100*(total number of symptoms triggered on an object/total number of symptoms)).

The compliance score for the object is based on the most critical of the violated standards. The score card displays 100 when all objects are compliant. When an object is non-compliant, the number of non-compliant symptoms are displayed in red and the total number of symptoms in grey.

> **NOTE**
> The compliance score for a user with limited object visibility is the same as for a user with complete object visibility. This is because the compliance score is calculated for all objects, irrespective of whether the user has access to the object or not.

#### Where You Find Compliance Score Cards

You can view score cards for each of the different types of benchmarks in the Compliance page available under **Operations › Compliance**.

You can view score cards for objects in the **Inventory Page › Compliance** tab.

#### Compliance Page

In the **Operations › Compliance** summary page, VMware Aria OperationsVMware Cloud Foundation Operations monitors compliance for VMware SDDC benchmarks.

VMware Aria OperationsVMware Cloud Foundation Operations displays compliance score cards for the following cards:

- vCenter
- VMware Cloud Foundation

- Google Cloud VMware Engine

Score cards are displayed for:

- VMware SDDC Benchmarks
- Custom Benchmarks
- Regulatory Benchmarks

**Compliance Tab**

In the **Inventory › Compliance** tab, VMware Aria OperationsVMware Cloud Foundation Operations displays score cards for the benchmarks that include the current objects in their calculations, based on the alert definitions and policies associated with that benchmark. The score cards display the total number of rules and the number non-compliant (violated) rules based on symptoms for each hardening guide.

**Score Cards in the Compliance Page**

In the **Operations › Compliance** page, you can view scores for benchmarks that you have activated. Click a score card to view more information.

**Table 147: Compliance Page Score Card Options**

| Item | Description |
|---|---|
| Score card for the configured hardening guides, custom benchmark and management packs | Displays the compliance score, total compliant and non-compliant objects for the compliance standards you have configured. |
| Object Breakdown | Displays the number of compliant and non-compliant objects for the following types of objects:<br>• vCenter<br>• ESXi Host<br>• Virtual Machine<br>• Distributed Port Group<br>• Distributed Virtual Switch<br>• vSAN Cache Disk<br>• vSAN Capacity Disk<br>• vSAN Cluster |
| Compliance Alert List | A list of alerts, grouped by time by default. You can either remove the grouping of the alerts, or group by criticality, definition, and object type.<br><br>The alerts which caused the compliance violation are displayed in a table. You can sort the table by the following columns:<br>• Alert ID<br>• Criticality<br>• Alert<br>• Triggered On<br>• Updated On<br><br>Click **Actions** to edit the policies applied to the configuration guide.<br><br>Click an alert to view more details. The **Alerts** tab in the **Inventory** page opens. |

*Table continued on next page*

*Continued from previous page*

| Item | Description |
|---|---|
| | Use the advanced quick filter to search for alerts. Click the drop-down arrow for more options to narrow your search. |

### Compliance Alerts

You use the compliance score card as an investigative tool when you evaluate the state of objects in your environment, or when you research the root cause of a problem. If the score card indicates a problem, you can view the alerts to see details about the violation. Violated rules are based on the symptoms defined in the compliance alert.

The compliance alerts, which have the subtype named Compliance, include one or more symptoms that represent the compliance rules. Compliance alerts that are triggered appear in the **Compliance** tab in an object when you navigate from **Inventory** on the left menu. Compliance alerts appear as violations to the standard, and the triggered symptoms appear as violated rules. The rules are the alert symptoms, and the symptom configuration identifies the incorrect value or configuration. If a rule symptom is triggered for any of the alerts in the standard, the triggered rule violates the standard and affects the score that appears on the **Compliance** tab.

**Table 148: Compliance Tab Alert Display**

| Item | Description |
|---|---|
| Score card for the configured hardening guides | Displays the score card value, total number of rules, and number of non-compliance rules for the compliance standards you have configured. |
| Active Compliance Alerts | If you click the score card, the rules for the score card appear. When a symptom is triggered, the rule is considered to be violated. View the list of rules in the following tabs:<br>• Violated Rules. Displays only the triggered symptoms. Click a symptom to view more information.<br>• All Rules. Displays triggered and untriggered symptoms. |

## How To Configure Compliance Benchmarks

Configure VMware SDDC, custom, and regulatory benchmarks from the Compliance page. Unlike previous releases, you can now enable alert definitions in one of the active policies, from the Compliance page directly.

### Activate VMware Cloud Foundation Benchmarks

You can activate the VMware Cloud Foundation benchmarks to audit the compliance of the objects in the VMware Cloud Foundation stack. The products currently included for assessment are ESXi, SDDC Manager, vCenter Server, vSAN and NSX.

You must configure the VMware Cloud Foundation data source before you proceed. For information, see the topic, VMware Cloud Foundation in the *Configuring VMware Aria Operations* guide.
Before you can enable the VMware Aria Operations CSA Compliance Pack for VMware Cloud Foundation, you must install the integration after downloading the .PAK file from Marketplace. For more information, see the VMware Aria Operations for Integrations Documentation.

1. From the left menu, click **Operations › Compliance** to access the compliance page.
2. Select VMware Cloud Foundation from the carousel on the top of the page.

3. Do the following:
   a) To enable VCF Compliance based on the VCF Audit Guides, in the **VCF Benchmarks** section, click **Enable** in the **VCF Compliance based on the VCF Audit Guides** card.
   b) To enable VMware Aria Operations CSA Compliance Pack for VMware Cloud Foundation, in the **Regulatory Benchmarks** section, click **Enable** in the **VCF Compliance based on the Cloud Security Alliance** card.
4. Select the policy that you want to modify. When there are child policies, you can select a child policy and unselect a parent policy. VMware Aria Operations modifies the selected policy and activates the alert definitions associated with the current scorecard.
5. Click **Enable** to confirm your selection.

VMware Aria Operations starts to assess the objects based on the policy that you selected. To edit a policy, click **Edit** in the configuration guide pane and select a different policy.

### Activate VMware SDDC Benchmarks

You can activate the VMware SDDC Benchmark to monitor objects for violation of vSphere Security Configuration Guide, vSAN Security Configuration Guide, NSX Security Configuration Guide (SDDC only). The score cards in the VMware SDDC Benchmark warn you when compliance alerts trigger on your vCenter Server instance, NSX-V objects, NSX objects, vSAN objects, ESXi hosts, virtual machines, distributed port groups, or distributed virtual switches.

You must configure the data source for which you are enabling the VMware SDDC benchmarks before you proceed. For information on how to configure a data source, see the relevant topic in the Integrating Data Sources with VMware Aria OperationsVMware Cloud Foundation Operations chapter of the *Configuring VMware Aria Operations* guide.

1. From the left menu, click **Operations › Compliance** to access the compliance page.
2. Select a data source from the carousel on the top of the page.
3. In the VMware SDDC Benchmarks section, click **Enable** under the vSphere Security Configuration Guide, vSAN Security Configuration Guide or NSX Security Configuration Guide card.
   The **Enable Policies** dialog box opens.
4. Select the policy that you want to modify. When there are child policies, you can select a child policy and unselect a parent policy. VMware Aria OperationsVMware Cloud Foundation Operations modifies the selected policy and activates the alert definitions associated with the current scorecard.
5. Click **Enable** to confirm your selection.

VMware Aria OperationsVMware Cloud Foundation Operations starts to assess the objects based on the policy that you selected. To edit a policy, click **Edit** in the configuration guide pane and select a different policy.

### Create a New Custom Benchmark

You can create a custom compliance benchmark to ensure that objects comply with compliance alerts available in VMware Aria OperationsVMware Cloud Foundation Operations, or custom compliance alert definitions. When a compliance alert is triggered on your vCenter instance, hosts, virtual machines, distributed port groups, or distributed switches, you investigate the compliance violation. You can add up to five custom compliance score cards.

To create a custom benchmark based on industry standard regulatory compliance requirements, you must first download and install the compliance management packs.

1. From the left menu, click **Operations › Compliance** to access the compliance page.

2. In the Custom Benchmarks section, click **Add Custom Compliance**.
   The **Add Custom Compliance** dialog box opens.
3. Select **Create a New Custom Benchmark**.
   a) In the Name and Description step, provide a name and description for the custom benchmark and click **Next**.
   b) In the Alert Definitions step, select the compliance alerts that you want to add to this custom compliance benchmark and click **Next**.
   c) In the Policies step, select the policies to activate compliance and click **Finish**.

The custom compliance which monitors alert definitions that you selected is available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit**.

### Import or Export a Custom Benchmark

You can export custom benchmarks from any VMware Aria OperationsVMware Cloud Foundation Operations instance and import it to another instance. Reusing custom benchmarks saves you time and effort. You can modify an imported custom benchmark. Exported files are in the XML format. The XML file contains information about alert groups, alerts, and filters.

You must first export a XML file with the custom benchmarks from another instance of VMware Aria OperationsVMware Cloud Foundation Operations before importing the XML file to another instance.

1. From the left menu, click **Operations › Compliance** to access the compliance page.
2. To import a custom benchmark, select either the SDDC or the VMC SDDC tab depending on where your objects are present.
3. In the Custom Benchmarks section, click **Add Custom Compliance**.
   The **Add Custom Compliance** dialog box opens.
4. Select **Import An Existing Custom Benchmark**.
   a) In the Import Compliance Score card dialog box, select the Score card Definition XML file from your local computer. If the XML file contains cloned alerts from the VMware Aria OperationsVMware Cloud Foundation Operations instance that was used to export the file, the cloned alerts are also imported.
   b) VMware Aria OperationsVMware Cloud Foundation Operations displays a message to indicate if the XML file was successfully imported.
   c) If you see a message which indicates that there is a conflict between the data in the XML file and the custom benchmarks already defined, make a selection on how to handle a conflict.
   d) Click **Done**.
5. To export an existing custom benchmark, click the score card to select the benchmark and select **Export** from the **Actions** menu.

The imported compliance benchmarks are available in the Custom Benchmarks section of the Compliance page. You can edit the alert definitions and policies at any time by clicking **Edit** from the **Actions** menu after clicking the score card.

### Activate a Regulatory Benchmark

To enforce and report on the compliance of your vSphere objects, you activate the compliance pack that contains the policies for regulatory standards. Then, you select the policy to activate the appropriate regulatory alerts for your virtual machines.

Before you can enable some of the regulatory benchmarks, you may need to install the integration by following the steps in Installing and Configuring Integrations in VMware Aria Operations after downloading the .PAK file from Marketplace.

1. From the left menu, click **Operations › Compliance** to access the compliance page.
   The compliance packs for the regulatory standards are displayed under the Regulatory Benchmark section.
2. To activate any regulatory benchmark, click **Activate From Repository** on the required compliance pack.
   You are redirected to the **Management Pack Details** page.
3. Click **Activate** to complete the installation.
4. To activate the compliance pack policies, navigate to the **Compliance** homepage and click **Enable** on the installed compliance pack.
   The **Enable Policies** window opens.
5. Select the policies that you want to activate and click **Enable** to complete the process.

VMware Aria OperationsVMware Cloud Foundation Operations starts to assess the objects based on the regulatory benchmark that you installed.

## Viewing and Configuring Audit Events

You can examine the diverse log events generated across resources in a VCF instance, encompassing vSphere resources like vCenter Server, VMs, Hosts, vSAN resources, and NSX resources. By default, the Audit Events displays events from the past 24 hours, but you can change the display to view events in other time ranges. Each audit event has a detailed view containing information which helps you analyse the events in your virtual infrastructure for suspicious user actions or system issues.

## What are Audit Events?

VMware Aria Operations allows you to view, monitor and search through activities across all the vCenter Servers in your infrastructure.

Audit events enhances operational efficiency, transparency, and accountability. These improvements facilitate security and compliance audits and provide essential details for conducting security forensics. Your organization may be required to provide audit data to support both internal and external audit reviews. By viewing audit events, you can reduce the time required for audit cycles by collecting data from various resources across vCenter Servers, and see it in an unified interface. Currently, audit events reports on authentication and security-related audit events.

### How Does Audit Logging Work?

VMware Aria Operations can display audit events after integration with VMware Cloud Foundation Operations for logs. Audit events include interactions within the platform like searches, logins, logouts, capability checks, and configuration modifications, which result in the creation of relevant audit records.

By default, the VMware Aria Operations displays audit events for vCenter Server, VMs, Hosts, vSAN resources, NSX resources.

If your virtual infrastructure has multiple vCenter instances, VMware Aria Operations, can monitor resources across all of them. Audit events are only displayed for vCenter instances which are monitored by VMware Cloud Foundation Operations for logs , even if such instances are directly configured with VMware Aria Operations. You must manually configure VMware Cloud Foundation Operations for logs to monitor the vCenter instances in your virtual infrastructure. You can use log forwarding to configure multiple VMware Cloud Foundation Operations for logs instances to do the log analysis. Out of these instances, only one is integrated with VMware Aria Operations, while the others forward the logs to the integrated VMware Cloud Foundation Operations for logs instance.

**NOTE**
The instances of VMware Cloud Foundation Operations for logs used for log forwarding must be of the supported version. Otherwise, audit events are not displayed for vCenter instances monitored by those vCenter log forwarders.

For more information, see the following topics:
- Configure VMware Aria Operations for Logs to Pull Events, Tasks, and Alarms from vCenter Server Instance
- Using VMware Aria Operations with VMware Aria Operations for Logs
- Add a VMware Aria Operations for Logs Log Forwarding Destination

To summarize, VMware Aria Operations displays audit events from vCenter instances which are:
- Configured in the VMware Cloud Foundation Operations for logs instance which is integrated with VMware Aria Operations.
- Configured in the VMware Cloud Foundation Operations for logs instances which forward logs to the VMware Cloud Foundation Operations for logs which is integrated with VMware Aria Operations.

**Audit Event Categories**

VMware Aria Operations generates audit events for the following category of actions:
- Access
- Access_control
- Account_management
- Configuration
- Data_access
- Network
- Notification
- Permissions
- Session
- System
- Policy
- Firewall

Audit events may have more than one category.

The default time range for which the audit events are displayed is 24 hours. You have the option to select audit events from the past 1, 6, 12, 24 and 48 hours, but it depends on the data retention configured in the VMware Cloud Foundation Operations for logs instance. Additionally, you can specify a custom time range for any time period in the past, provided that it does not exceed the maximum limit of 48 hours.

## Configuring Audit Events

Configure integration with VMware Cloud Foundation Operations for logs to see audit events from vCenter Servers.

Ensure that you have the following versions:

- VMware Aria Operations version 8.18+
- VMware Cloud Foundation Operations for logs version 8.18+
  For more information, see the topic, Upgrading VMware Aria Operations for Logs, in the Aria Operations for Logs documentation.

Bi-directional integration between VMware Aria Operations and VMware Cloud Foundation Operations for logs is a prerequisite. This means, you must:
- Configure the VMware Cloud Foundation Operations for logs management pak within VMware Aria Operations, directing it to a VMware Cloud Foundation Operations for logs instance.

- Configure VMware Aria Operations integration within the VMware Cloud Foundation Operations for logs environment, pointing it to the VMware Aria Operations instance.

  **NOTE**
  Only one VMware Cloud Foundation Operations for logs instance can be integrated with VMware Aria Operations.

  **NOTE**
  You must manually configure VMware Cloud Foundation Operations for logs to monitor vCenter servers. Audit events can monitor and report audit events for vCenter Servers coming through I VMware Cloud Foundation Operations for logs instance which is connected via log forwarding.

1. From the left menu, click **Operations › Audit Events**.
2. If you have not configured VMware Aria Operations, you will not see any information related to audit events. Install VMware Cloud Foundation Operations for logs first.
   a) Install VMware Cloud Foundation Operations for logs.
   b) Register the vCenter Server instances in VMware Cloud Foundation Operations for logs.
   c) After you install VMware Cloud Foundation Operations for logs, you must configure the VMware Cloud Foundation Operations for logs management pack from the **Administration › Integrations** page in VMware Aria Operations.
   d) Configure the VMware Cloud Foundation Operations for logs instances which forward logs to the VMware Cloud Foundation Operations for logs which is integrated with VMware Aria Operations.
3. Verify that you can see the audit events in the **Audit Events** page.

## Viewing Audit Events

VMware Aria Operations displays audit events to provide comprehensive visibility into vCenter server activities, enabling monitoring and analysis of virtual infrastructure events for potential issues or security threats.

### Where You Find Audit Events

From the left menu, click **Operations › Audit Events** to access the Audit Events page. A list of audit events is displayed in a table in the right pane on the page. You can filter the audit events by using the options on the left pane. Click an event in the right pane to see more details. Use the search bar above the table to search for audit events. On the top of the page, you see the number of vCenter servers monitored.

**NOTE**
The audit events displayed on the page do not refresh automatically, when you are on the page. Use the **Reload** option to fetch the latest audit events.

### Changing the Time Range

By default, VMware Aria Operations displays audit events for the last 24 hours. You can change this time range by clicking the drop-down menu. The presets available are:

- Latest 1 hour
- Latest 6 hours
- Latest 12 hours
- Latest 24 hours
- Latest 48 hours

If you want to select a custom range, use the Start Date, End Date, Start Time and End Time options. The custom range is limited to a 48 hour period. Click **Apply** to filter the audit events.

**View Details**

You can click the **VIEW DETAILS** button to view information about the vCenter servers the audit events are logged from.

The **vCenter Servers Status** dialog box opens when you click the **VIEW DETAILS** button. Use the search bar to search for the monitoring status of specific vCenter servers. If a vCenter Server instance that you are aware is present in your datacenter is not being monitored, see the troubleshooting steps in the topic, Troubleshooting Audit Events.

## Troubleshooting Audit Events

There are three reasons why audit events for all your vCenter Servers are not displayed. Follow the troubleshooting steps if you do not see audit events from all the vCenter Servers in your infrastructure.

**Reason 1: vCenter Servers are registered with VMware Aria Operations but not with VMware Cloud Foundation Operations for logs.**

Solution: You must configure VMware Aria Operations for Logs to work with vCenter Server instances by performing the steps in the topic, Configure VMware Aria Operations for Logs to Pull Events, Tasks, and Alarms from vCenter Server Instance.

**Reason 2: VMware Cloud Foundation Operations for logs Log Forwarding instances are not of the required version.**

Solution: All versions of VMware Cloud Foundation Operations for logs must be of the 8.18+ version. You can perform the upgrade by following the steps in the topic, Upgrading VMware Aria Operations for Logs.

**Reason 3: Bidirectional sync between VMware Cloud Foundation Operations for logs and VMware Aria Operations has failed.**

Solution: You must configure the VMware Cloud Foundation Operations for logs management pack from the **Administration › Integrations** page in VMware Aria Operations. This ensures bidirectional sync. For more information, see the topic, Configuring VMware Aria Operations for Logs with VMware Aria Operations.

## Configuring Green Score to Track Sustainability

VMware Green Score is used to measure sustainability in digital operations through data center virtualization. Configure Green Score in VMware Aria Operations to track where you are in your sustainability journey and how it is progressing over time.

You can configure green score at the Organization level and for physical data centers.

## Configuring Green Score Sustainability Data at the Organization Level

Virtualization plays a crucial role in data center consolidation and thereby helps in the reduction of hardware footprint in the data centers. Virtualization helps in the saving of power consumption and floor space in data centers which improves the overall efficiency and curbs carbon emissions that result from IT infrastructure growth. Use this page to configure green score data in VMware Aria Operations. Using green score, organizations can track where they are in their sustainability journey and how they are progressing over time.

**NOTE**

Ensure that the following permissions are assigned:

- Home > Change Organization Settings
- Home > View Sustainability Page

**Configuring Green Score for the First Time**

From the **Home** page, click **Launchpad › Sustainability**. From the **Sustainability** page, click **Configure** from the Green Score panel. Enter the details in the **Organization** tab.

**Table 149: Organization Details**

| Option | Description |
|---|---|
| **Green Score Component Input - Virtualization Ratio(%)** | |
| Compute Virtualization | Enter the percentage of compute services that are virtual. Default value is 80%. The default values are based on the VMware estimate of the most common numbers. |
| Network Virtualization | Enter the percentage of network services that are virtual. Default value is 30%. The default values are based on the VMware estimate of the most common numbers. |
| Storage Virtualization | Enter the percentage of storage technologies that are virtual. Default value is 40%. The default values are based on the VMware estimate of the most common numbers. |
| Desktop Virtualization | Enter the percentage of desktops that are virtual. Default value is 20%. The default values are based on the VMware estimate of the most common numbers. |
| **Power Sources (%)** | |
| Add Power Source | Click this option to add the share of power sources that you use. The total should add up to 100%.<br>The power sources are coal, oil, natural gas, hydroelectric, solar PV, wind, biomass, and nuclear.<br><br>The default value for coal is 79.2% and the default value for hydroelectric is 20.8%. The default values are based on the global averages from the Greenhouse Gas (GHG) Emissions Intensity of Electricity Generation methods.<br><br>The actual global average value is 709g/kWh. |
| **Average Age of Hardware in years (as of today)** | |
| Servers (years) | Enter the average age of the servers in your organization. Default age is two years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |
| Storage (years) | Enter the average age of storage in your organization. Default age is two years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |
| Network (years) | Enter the average age of network in your organization. Default age is three years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |
| Desktop (years) | Enter the average age of the desktops in your organization. Default age is five years. The value changes dynamically each day. The default values are based on the VMware estimate of the most common numbers. |

After you click **Save**, the Sustainability data is calculated. Calculation of the sustainability data may take some time.

**Green Score Details**

Green Score is calculated based on five components, which are: Workload Efficiency, Utilization of Physical resources, the extent of Virtualization, Power Source, and Hardware Efficiency. Each of these components has varying levels of

impact on the carbon emissions within a data center and hence has a different weightage when contributing to a common Green Score.

- **Workload Efficiency**: Contributes to 22.5% weightage Efficiency = Wastage / (Wastage + Non Wastage).
- **Resource Utilization**: Contributes to 12.5% weightage based on the utilization of hardware from the absolute capacity without considering buffers and HA.
- **Virtualization**: Contributes to 15% weightage.
- **Power Source**: Contributes to 37.5% weightage. Scoring is based on the carbon intensity of the power source used.
- **Hardware Efficiency**: Contributes to 12.5% weightage, based on the fact that newer generation hardware will be more energy efficient.

**Table 150: Green Score Details and Dependencies**

| Name | Description |
|---|---|
| Greenscore Meter | Indicates how sustainable a company is. A higher score indicates a greater green score. Green score is calculated based on the different values calculated for workload efficiency, resource utilization, virtualization, power source, and hardware efficiency. |
| Greenscore Trend | Indicates the green score trend of the company. You can view daily, weekly, and monthly trends. |
| **Clean Demand** | |
| Workload Efficiency | Measures the wastage in the data center and calculates a score based on various waste sources. The workload efficiency value is calculated based on different metrics, such as powered off VMs, oversized VMs, idle VMs, snapshots, and orphaned disks. The metrics are as follows:<br>Total Consumption CPU = (vCenter) CPU\|vCPUs Allocated on all Powered On Consumers<br><br>Total Consumption Memory (GB) = (vCenter) Memory\|Host Usage (KB)<br><br>Total Consumption Disk (TB) = (vCenter) Disk Space\|Total Provisioned Disk Space (GB)<br><br>Oversized VM\|vCPU = (VM) Summary\|Oversized\|Virtual CPUs and Oversized VM\|Memory (GB) = (VM) Summary\|Oversized\|Memory (KB) - These metrics are used to calculate workload efficiency only for powered on VMs. This means that vCPUs and memory of powered off and oversized VMs are excluded from the calculation.<br><br>Idle VM\|Disk (TB) = (Cluster) Reclaimable\|Idle VMs\|Disk Space (GB)<br><br>Idle VM\|vCPU = (Cluster) Reclaimable\|Idle VMs\|CPU (vCPUs)<br><br>Idle VM\|Memory (GB) = (Cluster) Reclaimable\|Idle VMs\|Memory (KB) |

*Table continued on next page*

*Continued from previous page*

| Name | Description |
|---|---|
| | Powered Off VM\|Disk (TB) = (Cluster) Reclaimable\| Powered Off VMs\|Disk Space (GB)<br><br>Snapshot\|Disk (TB) = (Cluster) Reclaimable\|VM Snapshots\|Disk Space (GB)<br><br>Orphaned VMDK\|Disk (TB) = (Datastore) Reclaimable\| Orphaned Disks\|Disk Space (GB) |
| **Lean Operations** | |
| Resource Utilization | Provides a value based on server and storage utilization. The score is based on the utilization of hardware from absolute capacity without considering buffers and HA. Here are the metrics used:<br>Total ESXi Hosts CPU Utilization (GHz) = (vCenter) CPU\| Usage (MHz)<br><br>Total ESXi Hosts CPU Capacity (GHz) = (vCenter) CPU\| Total Capacity (MHz)<br><br>Total ESXi Hosts Memory Utilization (TB) = (vCenter) Memory\|Host Usage (KB)<br><br>Total ESXi Hosts Memory Capacity (TB) = (vCenter) Memory\|Total Capacity (KB) Organization/Physical<br><br>Total datastores capacity = (vCenter) Disk Space\|Total Capacity (GB)<br><br>Total datastores utilization = (vCenter) Disk Space\| Utilization (GB)<br><br>Total RDM = Sum((VM) Virtual Disk \ Configured Size) |
| Virtualization | Virtualization score is calculated based on the virtualization percentage of compute, storage, network, and desktop entered in the **Organization Details** page. The formula is a weighted average of virtualization percentages with 40%, 30%, 20%, and 10% corresponding weights.<br><br>**NOTE**<br>At the organization level, green score is calculated based on physical data centers that are configured for green score. In addition, green score calculation at the organization level also includes vCenter instances that are not associated with any physical data center. |
| **Green Supply** | |
| Power Source | The score is based on the carbon intensity of the power source used. For example, traditional coal based power sources will have a carbon intensity rate of 888 g/kWh while wind based power source will have a carbon intensity of 26 only (Reference : WNA Report - Comparison of |

*Table continued on next page*

*Continued from previous page*

| Name | Description |
|---|---|
| | Lifecycle Greenhouse Gas Emissions of Various Electricity Generation Sources.<br>**Note**: A 100% badge score indicates zero carbon emission.<br><br>Each value is calculated based on the share of the power sources entered in the **Organization Details** page, after which the average value is calculated. The value for each of the power resources is calculated as follows:<br>Green Factor Shares (100%) Adjusted Score (Green Factor * Shares)<br><br>Coal = 17% (Green Factor)<br><br>Oil = 31% (Green Factor)<br><br>Natural Gas = 53% (Green Factor)<br><br>Solar PV = 92% (Green Factor)<br><br>Biomass = 96% (Green Factor)<br><br>Nuclear = 97% (Green Factor)<br><br>Hydroelectric = 98% (Green Factor)<br><br>Wind = 98% (Green Factor)<br><br>**NOTE**<br>Shares are user inputs in the **Organization Details** page.<br><br>**NOTE**<br>At the organization level, green score is calculated based on physical data centers that are configured for green score. In addition, green score calculation at the organization level also includes vCenter instances that are not associated with any physical data center. |
| Hardware Efficiency | The hardware efficiency badge is based on the fact that new generation hardware is more energy efficient than the old generation hardware - by running more workloads on the same sized hardware and also the ability of hardware to optimize energy consumption. The efficiency of hardware which includes server, storage, network, and desktop is based on the average age of the hardware entered in the **Organization Details** page. For example, if you had entered the age of your servers as 2 years, then the hardware efficiency for servers would be 80%, because the useful life duration is considered as 10 years. |

*Table continued on next page*

*Continued from previous page*

| Name | Description |
|---|---|
| | **NOTE**<br>At the organization level, green score is calculated based on physical data centers that are configured for green score. In addition, green score calculation at the organization level also includes vCenter instances that are not associated with any physical data center. |

**Power Consumption Chart**

The power consumption chart displays the cumulative power consumption for the last one month. It is calculated using the power usage metrics and is the sum of the last month's daily power consumption. When you click the Power Consumption option, you can also view the **Power Consumption** trend which can be daily, weekly, or monthly.

**Carbon Footprint Chart**

The carbon footprint chart displays the carbon emission for the last one month. When you click the Carbon Footprint option, you can also view the **Carbon Footprint** trend which can be daily, weekly, or monthly.

**Environmental Impact**

The environmental impact is depicted in terms of carbon emissions from the number of smart phones charged, as per the calculations based on Greenhouse Gas Equivalencies Calculator. For detailed info on this please refer to https://www.epa.gov/energy/greenhouse-gases-equivalencies-calculator-calculations-and-references#smartphones.
For information about Sustainability dashboards, see Sustainability Dashboards.

## Configuring Green Score Sustainability Data for Physical Data Centers

Virtualization plays a crucial role in data center consolidation and thereby helps in the reduction of hardware footprint in the data centers. Virtualization helps in the saving of power consumption and floor space in data centers which improves the overall efficiency and curbs carbon emissions that result from IT infrastructure growth. Use this page to configure green score data in VMware Aria OperationsVMware Cloud Foundation Operations for physical data centers that have been added to a vCenter Server account or to a VMware Cloud Foundation account and configured for green score. Using green score, organizations can track where they are in their sustainability journey and how they are progressing over time.

**NOTE**
Ensure that the following permissions are assigned:
- Home > Change Organization Settings
- Home > View Sustainability Page

**Configuring Green Score for the First Time for Physical Data Centers**

1. From the **Home** page, click **Launchpad › Sustainability**.
2. From the **Sustainability** page, click **Configure** from the Green Score panel. Enter the details in the **Organization** tab.

   **NOTE**
   You must configure sustainability data at the Organization level and only then will you be able to configure sustainability for physical data centers. For information about configuring sustainability at the organization level, see Configuring Green Score Sustainability Data at the Organization Level.

3. Click the **Save** button in the **Organization** tab to save the Organization level details and to activate the **Physical Data Centers** tab.
4. From the **Physical Data Centers** tab, click the **Add Physical Data Center** button to add a physical data center to a vCenter Server account or to a VMware Cloud Foundation account if you have not previously added a physical data center from the **Administration › Control Panel › Physical Data Center** page. For more information about adding physical data centers, see Adding Physical Data Centers in VMware Aria OperationsVMware Cloud Foundation Operations.
5. When you create a physical data center, the green score configuration page for that physical data center is displayed and you can fill in all the details.

   **NOTE**
   If you have already added a added a physical data center from the **Administration › Physical Data Center** page, a list of physical data centers appears on the **Physical Data Centers** tab. You can select each of the physical data centers from the vertical ellipses called **Add Greenscore Configuration**.

6. Select a physical data center from the Physical Data Centers map/list and then enter the details in the **Physical Data Centers** tab.

**Table 151: Physical Data Center Details**

| Option | Description |
|---|---|
| **Green Score Component Input - Virtualization Ratio(%)** | |
| Compute Virtualization | Enter the percentage of compute services that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| Network Virtualization | Enter the percentage of network services that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| Storage Virtualization | Enter the percentage of storage technologies that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| Desktop Virtualization | Enter the percentage of desktops that are virtual. Default values for the physical data center green score configuration are inherited from Organization level inputs. |
| **Power Sources (%)** | |
| Add Power Source | Click this option to add the share of power sources that you use. The total should add up to 100%.<br>The power sources are coal, oil, natural gas, hydroelectric, solar PV, wind, biomass, and nuclear.<br><br>The default value for coal is 79.2% and the default value for hydroelectric is 20.8%. The default values are based on the global averages from the Greenhouse Gas (GHG) Emissions Intensity of Electricity Generation methods.<br><br>The actual global average value is 709g/kWh. |
| **Average Age of Hardware in years (as of today)** | |
| Servers (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Storage (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |
| Network (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |
| Desktop (years) | Enter the average age of the servers for your physical data center. Default values for the physical data center green score configuration are inherited from Organization level inputs. The value changes dynamically each day. |

After you click **Save**, the Sustainability data is calculated. Calculation of the sustainability data may take some time.

**Green Score Details**

Green Score is calculated based on five components, which are: Workload Efficiency, Utilization of Physical resources, the extent of Virtualization, Power Source, and Hardware Efficiency. Each of these components has varying levels of impact on the carbon emissions within a data center and hence has a different weightage when contributing to a common Green Score.

- **Workload Efficiency**: Contributes to 22.5% weightage Efficiency = Wastage / (Wastage + Non Wastage).
- **Resource Utilization**: Contributes to 12.5% weightage based on the utilization of hardware from the absolute capacity without considering buffers and HA.
- **Virtualization**: Contributes to 15% weightage.
- **Power Source**: Contributes to 37.5% weightage. Scoring is based on the carbon intensity of the power source used.
- **Hardware Efficiency**: Contributes to 12.5% weightage, based on the fact that newer generation hardware will be more energy efficient.

**Table 152: Green Score Details and Dependencies**

| Name | Description |
|---|---|
| Physical Data Centers (map/list) | Displays the configured physical data centers in a map view or in a list. Select the physical data center for which you want to see green score details in a map view or from a list using the options at the top right corner.<br>To reset to the Organizational level, click the **X** button for the selected physical data center. |
| Physical Data Centers by Greenscore | Displays the number of all physical data centers. At the bottom of the panel, you can view the number physical data centers that are configured (green in color) and not configured (gray in color) for green score. For configured physical data centers, you can view the range of green score values applicable to the configured physical data center. |
| Greenscore Meter | Indicates how sustainable a company is. A higher score indicates a greater green score. Green score is calculated based on the different values calculated for workload |

*Table continued on next page*

*Continued from previous page*

| Name | Description |
|---|---|
| | efficiency, resource utilization, virtualization, power source, and hardware efficiency. |
| Greenscore Trend | Indicates the green score trend of the company. You can view daily, weekly, and monthly trends. |
| **Clean Demand** | |
| Workload Efficiency | Measures the wastage in the data center and calculates a score based on various waste sources. The workload efficiency value is calculated based on different metrics, such as powered off VMs, oversized VMs, idle VMs, snapshots, and orphaned disks. The metrics are as follows:<br>Total Consumption CPU = (vCenter) CPU\|vCPUs Allocated on all Powered On Consumers<br><br>Total Consumption Memory (GB) = (vCenter) Memory\|Host Usage (KB)<br><br>Total Consumption Disk (TB) = (vCenter) Disk Space\|Total Provisioned Disk Space (GB)<br><br>Oversized VM\|vCPU = (VM) Summary\|Oversized\|Virtual CPUs and Oversized VM\|Memory (GB) = (VM) Summary\|Oversized\|Memory (KB) - These metrics are used to calculate workload efficiency only for powered on VMs. This means that vCPUs and memory of powered off and oversized VMs are excluded from the calculation.<br><br>Idle VM\|Disk (TB) = (Cluster) Reclaimable\|Idle VMs\|Disk Space (GB)<br><br>Idle VM\|vCPU = (Cluster) Reclaimable\|Idle VMs\|CPU (vCPUs)<br><br>Idle VM\|Memory (GB) = (Cluster) Reclaimable\|Idle VMs\|Memory (KB)<br><br>Powered Off VM\|Disk (TB) = (Cluster) Reclaimable\|Powered Off VMs\|Disk Space (GB)<br><br>Snapshot\|Disk (TB) = (Cluster) Reclaimable\|VM Snapshots\|Disk Space (GB)<br><br>Orphaned VMDK\|Disk (TB) = (Data store) Reclaimable\|Orphaned Disks\|Disk Space (GB) |
| **Lean Operations** | |
| Resource Utilization | Provides a value based on server and storage utilization. The score is based on the utilization of hardware from absolute capacity without considering buffers and HA. Here are the metrics used:<br>Total ESXi Hosts CPU Utilization (GHz) = (vCenter) CPU\|Usage (MHz) |

*Table continued on next page*

*Continued from previous page*

| Name | Description |
|---|---|
| | Total ESXi Hosts CPU Capacity (GHz) = (vCenter) CPU\|Total Capacity (MHz)<br><br>Total ESXi Hosts Memory Utilization (TB) = (vCenter) Memory\|Host Usage (KB)<br><br>Total ESXi Hosts Memory Capacity (TB) = (vCenter) Memory\|Total Capacity (KB) Organization/Physical<br><br>Total datastores capacity = (vCenter) Disk Space\|Total Capacity (GB)<br><br>Total datastores utilization = (vCenter) Disk Space\|Utilization (GB)<br><br>Total RDM = Sum((VM) Virtual Disk \ Configured Size) |
| Virtualization | Virtualization score is calculated based on the virtualization percentage of compute, storage, network, and desktop entered in the corresponding physical data center's configuration. The formula is a weighted average of virtualization percentages with 40%, 30%, 20%, and 10% corresponding weights. |
| **Green Supply** | |
| Power Source | The score is based on the carbon intensity of the power source used. For example, traditional coal based power sources will have a carbon intensity rate of 888 g/kWh while wind based power source will have a carbon intensity of 26 only (Reference : WNA Report - Comparison of Lifecycle Greenhouse Gas Emissions of Various Electricity Generation Sources.<br>**Note**: A 100% badge score indicates zero carbon emission.<br><br>Each value is calculated based on the share of the power sources entered in the **Organization Details** page, after which the average value is calculated. The value for each of the power resources is calculated as follows:<br>Green Factor Shares (100%) Adjusted Score (Green Factor * Shares)<br><br>Coal = 17% (Green Factor)<br><br>Oil = 31% (Green Factor)<br><br>Natural Gas = 53% (Green Factor)<br><br>Solar PV = 92% (Green Factor)<br><br>Biomass = 96% (Green Factor)<br><br>Nuclear = 97% (Green Factor) |

*Table continued on next page*

*Continued from previous page*

| Name | Description |
|---|---|
| | Hydroelectric = 98% (Green Factor)<br><br>Wind = 98% (Green Factor)<br><br>**NOTE**<br>Shares are user inputs in the corresponding physical data center's page. |
| Hardware Efficiency | The hardware efficiency badge is based on the fact that new generation hardware is more energy efficient than the old generation hardware - by running more workloads on the same sized hardware and also the ability of hardware to optimize energy consumption. The efficiency of hardware which includes server, storage, network, and desktop is based on the average age of the hardware entered in the **Organization Details** page. However, you can change and input ages that are specific to the physical data center. For example, if you had entered the age of your servers as 2 years, then the hardware efficiency for servers would be 80%, because the useful life duration is considered as 10 years. |

**Power Consumption Chart**

The power consumption chart displays the cumulative power consumption for the last one month. It is calculated using the power usage metrics and is the sum of the last month's daily power consumption. When you click the Power Consumption option, you can also view the **Power Consumption** trend which can be daily, weekly, or monthly.

**Carbon Footprint Chart**

The carbon footprint chart displays the carbon emission for the last one month. When you click the Carbon Footprint option, you can also view the **Carbon Footprint** trend which can be daily, weekly, or monthly.

**Environmental Impact**

The environmental impact is depicted in terms of carbon emissions from the number of smart phones charged, as per the calculations based on Greenhouse Gas Equivalencies Calculator. For detailed info on this please refer to https://www.epa.gov/energy/greenhouse-gases-equivalencies-calculator-calculations-and-references#smartphones.
For information about Sustainability dashboards, see Sustainability Dashboards.

## Configuring Automation Jobs

You can automate jobs to perform certain actions as per a schedule. You can create and manage automation jobs from the Automation Central page. Scheduling of jobs allows you to perform actions without manual supervision. For example, you can automate jobs to run during a maintenance window, which could be outside of working hours.

## Automation Central

Automation Central is where you can create jobs to automate optimization actions which reclaim or rightsize VMs. You can also schedule additional actions such as powering on, powering off, and rebooting VMs and instances. Once you set up recurring jobs, you can track and obtain reports on them. You can customize jobs so that they only run based certain on parameters. For example, if you choose to delete a snapshot as an action, you can specify how old the snapshot must be

before it is deleted. You can also schedule jobs from the Reclaim and Rightsizing pages, where you configure the job in the context of a recommendation provided by VMware Aria Operations.

## Actions that you can Automate

You can automate the following actions using Automation Central:

**Reclaim**

- Delete old snapshots
- Delete idle VMs
- Poweroff Idle VMs
- Delete powered off VMs

**Rightsize**

- Downsize oversized VMs
- Scale-up undersized VMs

**Other**

Schedule additional actions such as powering on, powering off, and rebooting VMs and instances. The additional actions are based on the public cloud (GCP, AWS, Microsoft Azure) or vCenter adapter.

## Where you find Automation Central

Click **Operations › Automation Central** in the left pane.

## How Automation Central Works

In the Automation Central page you see a list of upcoming jobs and a calendar under the **Schedule** tab. The calendar displays all the jobs that are scheduled for the current month. You can move between months to see more scheduled jobs.

### View the Summary of Scheduled Jobs

When you click on a date in the calendar, you see a summary of the job. The summary displays the frequency of the job, the type of job, and if the job is activated or deactivated. You can click **Preview** to see more details about the job, or click **Deactivate All Recurrences** to deactivate the job. To edit the job, click the **Edit** link.

### View a Report of Jobs

View reclamation and rightsizing reports. The reclamation report displays graphical and numerical data on the total cost saving, CPUs reclaimed, memory reclaimed and storage reclaimed for different time periods.

The rightsizing reports displays graphical and numerical data on the CPUs downsized, memory downsized, CPU oversized, and memory upsized for different time periods.

### View Job History

You can also view the history of configured jobs which have run. Click the **History** tab above the calendar to see the job name, and job details in a tabular format. Use the search to search for a Job or VM. Click the drop-down in the search box to perform an advanced search.

The job history page only displays the status of jobs. For detailed information about the task, go to **Administration › Control Panel › Recent Tasks**. Failed recurring jobs are triggered during the next run.

If VM Power Off is not allowed, the action fails and is indicate as such in the logs.

### View Configured Jobs

The **Jobs** tab is where you see a list of configured jobs. For each job, clicking the ellipses icon brings up a menu from where you can edit, delete, clone or deactivate the job. If a job that you created is not visible in the list, search for it by entering the name in the search box. Alternatively, you can check by the job status by clicking the drop-down in the search box for advanced search options.

**Prerequisites to Run Actions**

Actions are run through VMware tools. vCenter privileges required to run each actions is documented in topic, Privileges Required for Configuring a vCenter Adapter Instance. Automation Central calls vCenter APIs with the credentials supplied to the vCenter adapter in VMware Aria OperationsVMware Cloud Foundation Operations.

Automation Central is available for users who have the advanced license and above. You must have the necessary permission to schedule and manage jobs in the **Automation Central** page which you can access from **Operations › Automation Central**. An administrator can manage these permissions in the Roles tab available in **Administration › Control Panel › Access Control**. Make sure that the **Manage Job Schedules** permission is selected in the **Assign Permissions** section, under Operations. This is available when you expand the **Automation Central** category.

To schedule a job of a specific type, you must have permission to run the corresponding action defined under **Administration › Global Inventory › Actions** in the **Roles** tab when you navigate to **Administration › Control Panel › Access Control**.

**Operational Actions** must be enabled in the vCenter cloud adapter instance to run actions from Automation Central.

To use alternate user credentials for actions, see the topic, Configure a vCenter Server Cloud Account in VMware Aria Operations.

**Troubleshooting Automation Jobs**

To see the logs for scheduled jobs, go to **Administration › Control Panel › Support Logs › Others** and look for the following logs:

- actionScheduler-.log
- actions-data-.log

Or, go to **Administration › Control Panel › Support Logs › Analytics** and search for action execution in analytics-.log.

## Create Job from Automation Central

Create a job to schedule an action to be performed automatically. You can select the type of action you want to perform, and then select the scope of the action. You can filter the scope based on attributes and metrics. Every action has a configuration option, which enables you to control the execution of the job based on conditions.

1. On the left menu, click **Operations › Automation Central**. In the Automation Central page, click **Add Job**.
   The **Create New Job** page opens. This page displays the following cards:
   - Reclaim
   - Rightsize
   - Billing . This card is available when you when you have VMware Cloud Director configured in VMware Aria OperationsVMware Cloud Foundation Operations. For more information, see the topic, Scheduling Bill Generation Using Automation Central.
   - Actions
2. Click on a card based on the type of job you want to create.
3. In the first step of the wizard, specify the following properties to create the action:

| Property | Description |
|---|---|
| Name | Specify a name for the job. This is displayed in the calendar. |
| Description | Provide a description for the job. |
| Action Configuration | Select this option for scheduling VM optimization actions. The choices are: |

*Table continued on next page*

*Continued from previous page*

| Property | Description |
|---|---|
| | **Reclaim**:<br><br>1. Delete old snapshots. When you select this option, you must specify at least one of the following:<br>– The snapshot age in days.<br>– The snapshot size in MB using the operators from the drop-down list.<br>– The snapshot name, using one of the operators from the drop-down list.<br><br>**NOTE**<br>The criteria for snapshot deletion is a combination of all the options you choose.<br><br>2. Delete idle VMs. When you select this option, you must specify the idle time duration in days.<br>3. Power off idle VMs. When you select this option, you must specify the idle time duration in days.<br>4. Delete powered off VMs. When you select this option, you must specify the power off duration in days.<br><br>**NOTE**<br>The time definition for powered off VMs, idle VMs, and snapshots is configured **Capacity › R eclaim ›** page. Click the **Settings** link on the page to edit or view the reclamation settings. Ensure that you have reviewed the settings in the Reclaimation page before configuring the job here.<br><br>**Rightsize**:<br><br>1. Downsize oversized VMs. When you select this option, you must select the rightsize resource type, and the downsizing limit amount.<br>2. Scale-up undersized VMs. When you select this option, you must select the rightsize resource type.<br><br>**NOTE**<br>VMware Aria OperationsVMware Cloud Foundation Operations does not check if the VM hot add/remove setting is enabled. If the VM power off is not allowed, then the action will fail. You will find the error in Recent tasks and in the logs. |

*Table continued on next page*

*Continued from previous page*

<table>
<tr><th>Property</th><th>Description</th></tr>
<tr><td></td><td>
Actions
<ul>
<li>Adapter Type: Select a public cloud or vCenter adapter from the drop-down list. Options are, vCenter, GCP, AWS, Microsoft Azure.</li>
<li>Object Type: Select the object from the drop-down list on which you want the action to be performed. Options are, Virtual Machine, Host System, Cluster Compute Resource.</li>
<li>Action: Start typing the name of the action, or select an action from the drop-down list. The options are as follows:</li>
</ul>
<table>
<tr><th>Adapter Type</th><th>Available Action</th></tr>
<tr><td>vCenter</td><td>• Power Off VM<br>• Power On VM<br>• Reboot Guest OS For VM<br>• Suspend VM</td></tr>
<tr><td>GCP</td><td>• Power Off CE Instance<br>• Power On CE Instance<br>• Reboot CE Instance</td></tr>
<tr><td>AWS</td><td>• Power EC2 Instance<br>• Power On EC2 Instance<br>• Reboot EC2 Instance</td></tr>
<tr><td>Microsoft Azure</td><td>• Power Off Azure Virtual Machine<br>• Power On Azure Virtual Machine<br>• Reboot Azure Virtual Machine</td></tr>
</table>
NOTE<br>If the number of resources for a job is ten or less, then the job runs at once. If the number of resources is more than ten, then the jobs run in groups of ten, in parallel.
</td></tr>
</table>

4. Click **Next**.
5. In the second step of the wizard, select scope of the object which contains the VMs on which VMware Aria OperationsVMware Cloud Foundation Operations must run the automation jobs.

**NOTE**
- When you select scope for **Additional Actions**, the objects that you drag and drop must be either of the same object type, or the ancestor object (that will serve as a container) of the object type you have selected in the previous step. For example, if you selected **Host System** in the previous step, you need to select either a Host System object or an object that is higher in hierarchy (Cluster Compute Resources, Datacenter etc).
- VMware Aria OperationsVMware Cloud Foundation Operations ignores object types which do not match the object type that you selected in the previous step. Invalid object types can be object of descendant object types, objects from a different management pack/solution, or from unrelated object hierarchy.

a) Drag and drop the object type to the left pane.
b) To delete members, click the X icon.

6. After you have dragged & dropped the objects for the action, click **Preview Scope** at the bottom right corner to see on which objects the jobs will potentially be executed. This list may be different after you apply some filter from the **Filter Criteria** page.

   **NOTE**
   The final objects list on which the actions will be run can be seen by clicking the **Preview Scope** button for the relevant job in the **Schedule** tab.

7. Click **Next**.
8. In the **Filter Criteria** step of the wizard, you can create rules using which VMware Aria OperationsVMware Cloud Foundation Operations filters objects on which the jobs are run. If you do not create a rule, all the VMs within the scope of the selected object type are affected.
   a) In the **Set Filter Criteria** section, create rules using a combination of metrics, relationship, property, object name or tag values and logical expressions.
   b) You can add more than one criteria by clicking either the ADD ANOTHER CRITERIA SET or ADD buttons. When you click ADD ANOTHER CRITERIA SET, VMware Aria OperationsVMware Cloud Foundation Operations creates an OR logical expression between the criteria sets. When you click ADD button, VMware Aria OperationsVMware Cloud Foundation Operations creates an AND logical expression between the criteria.
9. Click **Next**.
10. In the **Schedule** step of the wizard, configure a schedule for the action to run.

| Property | Description |
|---|---|
| Start Date | Set a start date which is on, or after the current date, for the job to start. Select a date from the calendar displayed on the page.<br>**NOTE**<br>There will be a delay of up to five minutes for a job to run. The actual start time for an action also depends on the action itself and the number of target objects involved. |
| Time zone | For the start time and date to be calculated for the job, select a time zone. Choices are:<br>• Browser - based on the current location reported by the browser.<br>• Host - based on the current location of the host machine. |

*Table continued on next page*

*Continued from previous page*

| Property | Description |
|---|---|
| | • GMT based time zone - Based on a location calculated as per UTC+0. |
| Start Time | Select a start time for the job. The drop-down menu provides options in five minute intervals, starting closest to your current browser reported time. |
| Recurrence | Set a recurrence for the job. The choices are:<br>• **One-Time**. No further configuration is possible.<br>• **Daily**. Set the recurrence in days. Set how frequently it must run by providing a value in **Run Every** field.<br>• **Weekly**. Select the days of the week when you want the job to run, by clicking the day of the week abbreviation.<br>• **Monthly**. Select the months when you want the job to run, by clicking the month abbreviation. By default, all the months are selected. For the months that you select, you can configure the job to run:<br>– On specific days of the month, by clicking the number. Or, the last day of the month, without specifying the exact date, by clicking **Last**.<br>– On the specific number of the week (first, second, third, fourth or the last) in the month, combined with the specific day of the week.<br>You can make multiple selections for each of the options in the drop down menu.<br>For the **Daily**, **Weekly**, **Monthly** options, you can set when the job must end based on a date, or the number of occurrences. |

a) In the **Notifications** section, select the **Receive Updates on Job via Email** check box to receive notifications two hours before the job is set to run and after it has been executed. For the email to be sent, you must also select the email outbound plugin from the drop down menu, and enter the email address to which the email must be sent. If you have not created an email outbound plugin, see the topic, Outbound Settings in .

**NOTE**
You can send updates via email to only one email ID.

11. Click **Create** to complete the steps in the wizard and create the job.

## Create Job from Reclaim or Rightsizing

You can create an automation job based on the recommendation provided by VMware Aria OperationsVMware Cloud Foundation Operations in the Reclaim or Rightsizing pages. You cannot create an automation job outside of the context provided by VMware Aria OperationsVMware Cloud Foundation Operations here. Use Automation Central for that.

1. Do one of the following:
   a) In the menu, click **Home**, and then in the left pane, click **Optimize Capacity › Reclaim**.
   b) In the menu, click **Home**, and then in the left pane, click **Optimize Performance › Rightsizing**.

2. In the Reclaim or Rightsizing pages, do the following
   a) Click the data center that you want to optimize.
   b) In the table heading that displays, Select the types of VMs that you want to optimize.
   c) Click the name of a listed cluster to show its VM list.
   d) Select the checkbox next to the VM that you want to optimize.
   e) Click **SCHEDULE ACTION**.
3. In the **Create Schedule Job** dialog box that opens, configure the following parameters:

| Property | Description |
|---|---|
| Job Name | Provide a name for the job. This information is displayed in the calendar in the Automation Central page. |
| Job Description | Provide a description for the job. |
| Start Date | From the date picker, select a date when the automation job should start. |
| Time of Day | • From the time picker, select the start time for the job.<br>• From the drop down, select the time zone when the time that you selected is valid for. |
| Receive updates on Job via Email | Select this checkbox if you have an email server configured, and you want to receive email notification about the status of the job. Notifications are sent two hours before the execution of the job. |
| Notification Method | If you selected the previous option, select the email outbound plugin from the drop down menu, and enter the email address to which the email must be sent. |

4. Click **Create**.

The automation job is created, and is available in the Automation Central page. From there, you can preview, edit, or delete the job.

**NOTE**
The scope is not editable since the objects on which the job should be run have been already selected from the Reclaim/Rightsize page.

.

## Log Analysis with VMware Cloud Foundation Operations for logs

When VMware Aria OperationsVMware Cloud Foundation Operations is integrated with VMware Cloud Foundation Operations for logs, you can search for log messages in context, and collect and analyze log feeds from within VMware Cloud Foundation Operations. You can view log-related metrics for troubleshooting. You can also dynamically extract fields from log messages based on customized queries.

## Log Analysis

You can analyse logs in the following ways:
- From the left the menu, select **Operations › Log Analysis**.
- By selecting an inventory object from the **Global Inventory** page, and clicking the **Logs** tab.

You can also create a dashboard which displays the log analysis screen.

For more details on how to integrate VMware Cloud Foundation Operations with VMware Cloud Foundation Operations for logs, see the topic, Configuring VMware Cloud Foundation Operations for logs with VMware Aria Operations.

## The Log Analysis Page

On the Log Analysis page, you can quickly search for logs or filter down the logs displayed using the following options:

**Text Based Search**

- Search bar: On the top of the screen is a search bar to search for logs.
- Time duration of logs: Change this value to display logs from a different time duration in the botton of the screen. The default is 5 minutes. Use **Custom** for a custom time duration.

**Advanced filtering**

Use the advance filtering to narrow down the logs you are looking for. VMware Cloud Foundation Operations parses the logs and gives you advanced filtering capabilities.

- Filters: Use source (Use _li_source_path), hostname, severity, priority, and appname filters.
- Condition operators: Use filters with the contains, doesnotContain, startsWith, doesNotStartWith, matchesRegex, exist, and doesNotExist conditions.
- Query box: Input a query in the query box to filter the logs displayed.

The log analysis screen displays the event timeline in a bar chart by default. This tells you how the logs have spiked for the given time duration. You can correlate spikes in logs with metrics to narrow down the root cause of an issue. On the screen, you can change the event timeline frequency. You can switch from the bar graph to a line graph.

The logs are displayed in the bottom half of the page in a table which contains a timestamp and log column. In the table which displays the logs, you can expand a log to see more details.

## Viewing Logs in VMware Aria Operations for Logs

Click the **LAUNCH OPERATIONS FOR LOGS** button to view the VMware Cloud Foundation Operations logs in VMware Aria Operations for Logs.

## Configuring VMware Cloud Foundation Operations for logs with VMware Aria Operations

To use the Log Analysis page, the Troubleshoot with Logs dashboard, and **Logs** tab in VMware Aria Operations, you must configure VMware Cloud Foundation Operations for logs with VMware Aria Operations.

### Configuring the VMware Cloud Foundation Operations for logs Adapter in VMware Aria Operations

You can integrate only one VMware Cloud Foundation Operations for logs instance. For Service Roles in Logs, see the Requirements for Integrating With VMware Aria Operations.

**Prerequisites**

- Verify that VMware Cloud Foundation Operations for logs and VMware Aria Operations are installed.
- Verify that you know the Credentials and IP address or FQDN of the VMware Cloud Foundation Operations for logs instance you have installed.

**Procedure**

1. From the left menu, select **Administration › Integrations**.
2. From the **Accounts** tab in the **Integrations** page, click **Add Account** and then select the VMware VMware Cloud Foundation Operations for logs card. You can also activate the management pack from the **Repository** tab, where you will find the card in the **Available Integrations** section. After the management pack is activated, click **Add Account**.
3. In the VMware Cloud Foundation Operations for logs page, add details to the following properties:
   - Name: Provide a name for the VMware Cloud Foundation Operations for logs adapter
   - Description: Provide a description for the VMware Cloud Foundation Operations for logs adapter.

- Operations for Logs server: Enter the IP address or FQDN in the **Log Insight server** text box of the VMware Cloud Foundation Operations for logs you have installed and want to integrate with.
- Credential: Click the + icon to enter the credentials of the VMware Cloud Foundation Operations for logs account you want to integrate with. You must enter the **Credential name**, **Operations for Logs User Name** and **Operations for Logs Password**.

  **NOTE**
  If you are upgrading from vRealize Operations Manager 8.x to VMware Aria Operations, you must re-configure these details after the upgrade is complete.

  **NOTE**
  The following types of credentials are allowed when configuring the VMware Cloud Foundation Operations for logs adapter in VMware Aria Operations:
  - Local VMware Cloud Foundation Operations for logs user. The minimum required privileges for the read-only user are **Explore Logs › Explore Logs** and **Explore Logs › Extracted Fields**, both with **View Access**.
  - Workspace One Access user (formerly known as VMware Identity Manager).
  - AD (Active Directory) user.

- Collector/Group: Select the collector group from the **Collectors/Groups** drop-down menu.

4. Click **Test Connection** to verify that the connection is successful.
5. Click **Save**.

## Configuring VMware Aria Operations in VMware Cloud Foundation Operations for logs

You configure VMware Aria Operations in VMware Cloud Foundation Operations for logs in the following scenarios:

- To have alerts and metrics generated from VMware Cloud Foundation Operations for logs.

**Prerequisites**

- Verify that VMware Cloud Foundation Operations for logs and VMware Aria Operations are installed.
- Verify that you know the IP address, hostname, and password of the VMware Aria Operations instance you want to integrate with.
- Verify that the VMware Cloud Foundation Operations for logs management pack is activated in VMware Aria Operations.

**Procedure**

1. From the Administration page of VMware Cloud Foundation Operations for logs, click VMware Aria Operations under **Integrations** in the left pane. You see the VMware Aria Operations Integration page.
2. In the **Hostname** text box, enter the IP address or FQDN of the VMware Aria Operations instance you want to integrate with.

   **NOTE**
   If you are using a load balancer, use its IP address or FQDN as a hostname value.

3. In the **Username** and **Password** text boxes, enter the user name and password of the VMware Aria Operations instance you want to integrate with.
4. Select the relevant check boxes according to your preference:
   - To send alerts to VMware Aria Operations, select **Enable alerts integration**.
   - To let VMware Cloud Foundation Operations for logs to navigate from a log to an object in VMware Aria Operations, select **Enable launch in context**.
   - To calculate and send metrics to VMware Aria Operations, select **Enable metric calculation**.
5. Click **Test Connection** to verify that the connection is successful and accept the certificate if it is untrusted.
6. Click **Save**.
   You can now view the log details for an object in VMware Aria Operations.

## Logs Tab

When VMware Aria Operations is integrated with VMware Cloud Foundation Operations for logs, you can view the logs for a selected object from the Logs tab. You can troubleshoot a problem in your environment by correlating the information in the logs with the metrics. You can then most likely determine the root cause of the problem.

**Where You Find the Logs Tab**

In the left menu, select an select an object from the **Global Inventory** in the left menu. Click the **Logs** tab. To view the **Logs** tab, you have to configure VMware Aria Operations in VMware Cloud Foundation Operations for logs. For more information, see the topic, Configuring VMware Cloud Foundation Operations for logs with VMware Aria Operations.

After integrating VMware Aria Operations with VMware Cloud Foundation Operations for logs, refresh the browser to see the **Logs** tab. For more information on how to use the Log Analysis screen, see the topic, Log Analysis with ,

## Viewing Logs in VMware Cloud Foundation Operations for logs

When VMware Aria Operations is integrated with VMware Cloud Foundation Operations for logs, you can search and filter log events. From the Interactive Analytics tab in the Logs page, you can create queries to extract events based on timestamp, text, source, and fields in log events . VMware Cloud Foundation Operations for logs presents charts of the query results.

To access the Logs page from VMware Aria Operations, you must either:

- Configure the VMware Cloud Foundation Operations for logs adapter from the VMware Aria Operations interface, or
- Configure VMware Aria Operations in VMware Cloud Foundation Operations for logs.
  For more information about configuring, see Configuring VMware Cloud Foundation Operations for logs with VMware Aria Operations.

For information about VMware Cloud Foundation Operations for logs interactive analytics, see the VMware Cloud Foundation Operations for logs documentation.

## Log Forwarding

For troubleshooting in the product UI, you can send the logs to an external log server or a VMware Cloud Foundation Operations for logs server.

If you have configured log forwarding from **Administration › Log Forwarding** in earlier versions of VMware Aria Operations, VMware recommends that you reconfigure in this version of VMware Aria Operations.

**Where You Find the Log Forwarding Page**

In the left menu, configure log forwarding from **Log Forwarding** tile under **Administration › Control Panel**.

**Table 153: Log Forwarding Page Options**

| Options | Description |
|---|---|
| Self-monitoring logging configuration | Forwards the logs to an external log server. |
| Forwarded Logs | You can select the set of logs you want to forward to the external log server or the VMware Cloud Foundation Operations for logs server. |
| Log Insight Servers | You can select an available VMware Cloud Foundation Operations for logs server IP. |

*Table continued on next page*

*Continued from previous page*

<table>
<tr><th>Options</th><th>Description</th></tr>
<tr><td></td><td>If there is no available VMware Cloud Foundation Operations for logs server IP, select Other from the drop-down menu and manually enter the configuration details.</td></tr>
<tr><td>Host</td><td>IP address of the external log server where logs have to be forwarded.</td></tr>
<tr><td>Protocol</td><td>You can select either<br><code>cfapi</code><br>or<br><code>syslog</code><br>from the drop-down menu to send event logging messages.</td></tr>
<tr><td>Port</td><td>The default port value depends on whether or not SSL has been set up for each protocol. The following are the possible default port values:<br><table><tr><th>Protocol</th><th>SSL</th><th>Default Port</th></tr><tr><td><code>cfapi</code></td><td>No</td><td>9000</td></tr><tr><td><code>cfapi</code></td><td>Yes</td><td>9543</td></tr><tr><td><code>syslog</code></td><td>No</td><td>514</td></tr><tr><td><code>syslog</code></td><td>Yes</td><td>6514</td></tr></table></td></tr>
<tr><td>Use SSL</td><td>Allows the VMware Cloud Foundation Operations for logs agent to send data securely.</td></tr>
<tr><td>Path to Certificate Authority File</td><td>You can enter the path to the trusted root certificates bundle file. If you do not enter a certificate path, the VMware Cloud Foundation Operations for logs Windows agent uses system root certificates and the VMware Cloud Foundation Operations for logs Linux agent attempts to load trusted certificates from <code>/etc/pki/tls/certs/ca-bundle.crt</code> or <code>/etc/ssl/certs/ca-certificates.crt</code>.</td></tr>
<tr><td>Cluster Name</td><td>Displays the name of the cluster. You can edit this field.</td></tr>
</table>

### Modifying Existing Log Types

If you manually modified the existing entries or logs sections and then modify the log forwarding settings from VMware Aria Operations, you lose the changes that you made.

The following server entries are overwritten by the VMware Aria Operations log forwarding settings.

```
port

proto

hostname

ssl

reconnect

ssl_ca_path
```

The following `[common | global]` tags are being added or overwritten by the VMware Aria Operations log forwarding settings.

`vmw_vr_ops_appname`

`vmw_vr_ops_clustername`

`vmw_vr_ops_clusterrole`

`vmw_vr_ops_hostname`

`vmw_vr_ops_nodename`

> **NOTE**
> Cluster role changes do not change the value of the `vmw_vr_ops_clusterrole` tag. You can either manually modify or ignore it.

## Configuring Alerts and Actions in VMware Aria OperationsVMware Cloud Foundation Operations

Configuring Alerts and Actions
Alerts are a way to be watchful of anything that is new or an issue that could be potentially dangerous to your environment. Whenever there is a problem in the environment, the alerts are generated. You can also create new alert definitions so that the generated alerts inform you about the problems in the monitored environment. In VMware Aria OperationsVMware Cloud Foundation Operations, alerts and actions play key roles in monitoring the objects.

You may have several questions while working on alerts. Following are some of the key questions that will help you navigate through the alert documentation.

- Where can I find all my alerts?
- What are the different types of alerts?
- Where can I get more information about alerts?
- How do I create a new alert definition?
- How do I define new symptoms?
- How do I define recommendations for my alerts?
- How can I create notification rules for alerts?
- How do I add my Outbound Plugins?
- How do I create payload templates for my Outbound Plugins?
- How can I export or import Outbound Settings?
- How do I deactivate alerts?
- How do I group alerts?
- What is intelligent alert clustering and how does it help in alert noise reduction?

Actions allow you to make changes to the objects in your environment. When you grant a user access to actions in VMware Aria OperationsVMware Cloud Foundation Operations, that user can take the granted action on any object that VMware Aria OperationsVMware Cloud Foundation Operations manages. For details, see Actions in VMware Aria Operations.

## Alert Definitions in VMware Aria OperationsVMware Cloud Foundation Operations

Alert Definitions

Alert definitions are a combination of symptoms and recommendations that you combine to identify problem areas in your environment and generate alerts on which you can act for those areas. You can then respond to the alerts with effective solutions that are provided in the recommendations.

## Where You Find Alert Definitions

To manage your alert definitions, from the left menu, click **Operations › Configurations**, and then click the **Alert Definitions** tile.

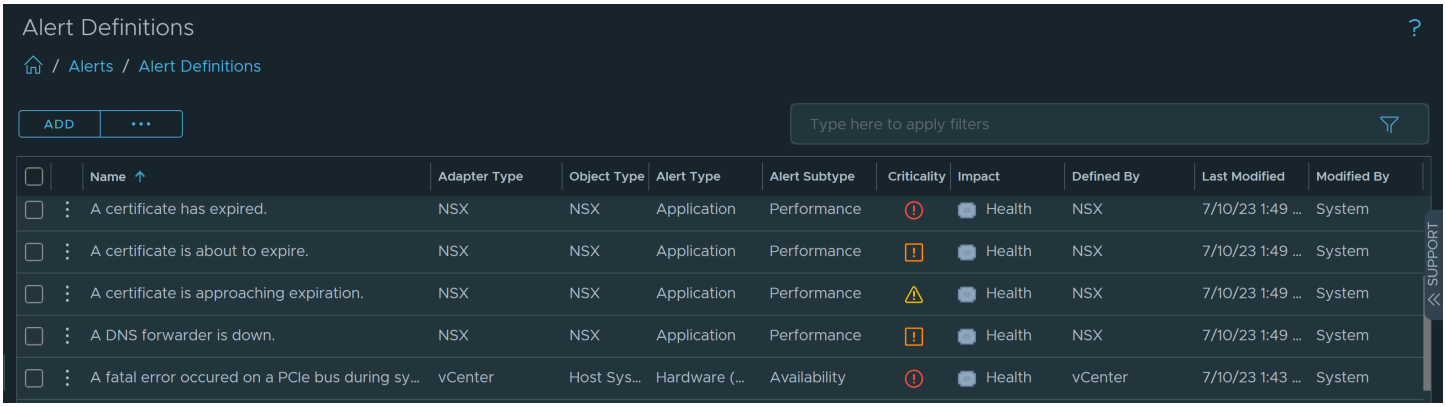


| Option | Description |
|---|---|
| Toolbar options | Use the toolbar options to manage your alert definitions.<br>• Add. Add an alert definition.<br><br>Click the horizontal ellipsis to perform the following actions.<br>• Edit. Modify the selected definition.<br>• Delete. Remove the selected definition.<br>• Clone. Create a copy of the selected definition so that you can customize it for your needs.<br>• Export. Downloads the alert definition.<br>• Import. Allows you to import alert definitions. To import:<br>– Click the **Import** option from the horizontal ellipsis.<br>– Click **Browse** and select the file to import.<br>– Select if you want to Overwrite or Skip the file in case of a conflict.<br>– Click **Import** to import the alert definition, and click **Done**. |
| Filters | Limits the list of alerts to those matching the filter you create.<br><br>You can also sort on the columns in the data grid. |
| Name | Name of the alert definition, which is also the name of the alert that appears when the symptoms are triggered. |
| Adapter Type | Adapter that manages the selected base object type. |
| Object Type | Base object type against which the alert is defined. |
| Alert Type | Metadata that is used to classify the alert when it is generated. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | You define the value on the Alert Impact page of the workspace. |
| Alert Subtype | Subcategory of the alert type and is the metadata that is used to classify the alert when it is generated.<br><br>You define the value on the Alert Impact page of the workspace. |
| Criticality | Severity of the alert when it is generated. The criticality includes the following possible values:<br>• Symptom. Alert is configured to display symptom based criticality.<br>• Critical<br>• Immediate<br>• Warning<br>• Info |
| Impact | Alert is configured to affect the Health, Risk, or Efficiency badge. |
| Defined by | Indicates who added the alert definition. The alert can be added by an adapter, a user, or the VMware Aria OperationsVMware Cloud Foundation Operations system. |
| Last Modified | Displays the date on which the alert was last modified. |

Predefined alerts are provided in VMware Aria OperationsVMware Cloud Foundation Operations as part of your configured adapters. Use Alert Definitions to manage your VMware Aria OperationsVMware Cloud Foundation Operations alert library, and to add or modify the definitions.

### Modifying Alert Definitions

If you modify the alert impact type of an alert definition, any alerts that are already generated will have the previous impact level. Any new alerts will be at the new impact level. If you want to reset all the generated alerts to the new level, cancel the old alerts. If they are generated after cancellation, they will have the new impact level.

### Symptoms in Alert Definitions

Symptom definitions evaluate conditions in your environment that, if the conditions become true, trigger a symptom and can result in a generated alert. You can add symptom definitions that are based on metrics or super metrics, properties, message events, fault events, or metric events. You can create a symptom definition as you create an alert definition or as an individual item in the appropriate symptom definition list.

When you add a symptom definition to an alert definition, it becomes a part of a symptom set. A symptom set is the combination of the defined symptom with the argument that determines when the symptom condition becomes true.

An alert definition comprises one or more symptom sets. If an alert definition requires all of the symptom sets to be triggered before generating an alert, and only one symptom set is triggered, an alert is not generated. If the alert definition requires only one of several symptom sets to be triggered, then the alert is generated even though the other symptom sets were not triggered.

### Recommendations in Alert Definitions

Recommendations are the remediation options that you provide to your users to resolve the problems that the generated alert indicates.

When you add an alert definition that indicates a problem with objects in your monitored environment, add a relevant recommendation. Recommendations can be instructions to your users, links to other information or

instruction sources, or VMware Aria OperationsVMware Cloud Foundation Operations actions that run on the target systems.

### Creating Alert Definitions

The alert definition process includes adding symptoms that trigger an alert and recommendations that help you resolve the alert. The alert definitions you create with this process are saved to your VMware Aria OperationsVMware Cloud Foundation Operations Alert Definition Overview list and actively evaluated in your environment based on your configured policies.

You can create or reuse existing symptoms and recommendations while defining an alert definition. If you create symptoms and recommendations, you add them to the definition, and they are added to the symptom and recommendations content libraries for future use. You also activate policies and select notifications for the alerts.

1. To create or edit your alert definitions, from the left menu, click **Operations › Configurations**, and then click the **Alert Definitions** tile.
2. Click **Add** to add a definition, or click the vertical ellipsis and select **Edit** to edit the selected definition.
3. In the **Alert** tab, enter details of the alert.

| Option | Description |
|---|---|
| Name | Name of the alert as it appears when the alert is generated. |
| Description | Description of the alert as it appears when the alert is generated. Provide a useful description for your users. |
| Base Object Type | The object type against which the alert definition is evaluated and the alert is generated.<br><br>The drop-down menu includes all of the object types in your environment. You can define an alert definition based on one object type. |
| Impact | Under Advanced Settings, select the badge that is affected if the alert is generated.<br><br>You can select a badge based on the urgency of the alert.<br>• Health. Alert requires immediate attention.<br>• Risk. Alert should be addressed soon after it is triggered, either in days or weeks.<br>• Efficiency. Alert should be addressed in the long term to optimize your environment. |
| Criticality | Severity of the alert that is communicated as part of the alert notification.<br><br>Select one of the following values.<br><br>• Info. Informational purposes only. Does not affect badge color.<br>• Warning. Lowest level. Displays yellow.<br>• Immediate. Medium level. Displays orange.<br>• Critical. Highest level. Displays red. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Symptom Based. In addition to alert criticality, each symptom includes a defined criticality. Criticality of the alert is determined by the most critical of all of the triggered symptoms. The color is dynamically determined accordingly. It you negate symptoms, the negative symptoms to not contribute to the criticality of a symptom-based alert. |
| Alert Type and Subtype | Select the type and subtype of alert.<br><br>This value is metadata that is used to classify the alert when it is generated, and the information is carried to the alert, including the alert notification.<br><br>You can use the type and subtype information to route the alert to the appropriate personnel and department in your organization. |
| Wait Cycle | The symptoms included in the alert definition remain triggered for this number of collection cycles before the alert is generated.<br><br>The value must be 1 or greater.<br><br>This setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition is added to the wait cycle for the symptom definitions. In most definitions you configure the sensitivity at the level of symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that after all of the symptoms are triggered at the desired symptom sensitivity level, the alert is immediately triggered. |
| Cancel Cycle | The symptoms are cancelled for this number of collection cycles after which the alert is cancelled.<br><br>The value must be 1 or greater.<br><br>This setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition is added to the cancel cycle for the symptom definitions. In most definitions you configure the sensitivity at the level of symptom level and configure the wait cycle of the alert definition to 1. This configuration ensures that after all of the symptom conditions disappear after the desired symptom cancel cycle, the alert is immediately canceled. |

4. Click **Next** to add symptom definitions.
5. In the **Symptoms/Conditions**, drag the selected symptom/condition in to the left pane. Use the workspace on the left to specify whether all or any of the symptoms/conditions or symptom/condition sets must be true to generate an alert. As you add one or more symptoms, you create a symptom expression. If this expression is evaluated as true,

then the alert is generated. You can similarly define one or more conditions for your alert, and when the conditions are met, the alert is generated. You can view the alert in the All Alerts page.

**Table 154: Add Symptoms/Conditions Selection Options**

| Option | Description |
|---|---|
| Defined On | Object that the symptom evaluates.<br><br>As you create alert definitions, you can select or define symptoms for the base object type and for related object types, based on the object relationship hierarchy. The following relationships are object types as they relate to the alert definition base object type.<br>• Self. A base object type for the alert definition. For example, host system.<br>• Descendant. An object type that is at any level below the base object type, either a direct or indirect child object. For example, a virtual machine is a descendant of a host system.<br>• Ancestor. An object type that is one or more levels higher than the base object type, either a direct or indirect parent. For example, a data center and a vCenter are ancestors of a host system.<br>• Parent. An object type that is in an immediately higher level in the hierarchy from the base object type. For example, a data center is a parent of a host system.<br>• Child. An object type that is one level below the base object type. For example, a virtual machine is a child of a host system. |
| Symptoms tab | |
| Select Symptom | Select the type of symptom definition that you are adding for the current Defined On object type.<br><br>• Metric / Property. Add symptoms that use metric and property symptoms. These metrics are based on the operational or performance values, and configuration properties that VMware Aria OperationsVMware Cloud Foundation Operations collects from target objects in your environment.<br>• Message Event. Add symptoms that use message event symptoms. These symptoms are based on events received as messages from a component of VMware Aria OperationsVMware Cloud Foundation Operations or from an external monitored system through the system's REST API.<br>• Fault Event. Add symptoms that use fault symptoms. These symptoms are based on events that monitored systems publish. VMware Aria OperationsVMware Cloud Foundation Operations correlates a subset of these events and delivers them as faults. Faults are intended to signify events |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | in the monitored systems that affect the availability of objects in your environment.<br>• Metric Event. Add symptoms that use metric event symptoms. These symptoms are based on events communicated from a monitored system where the selected metric violates a threshold in a specified manner. The external system manages the threshold, not VMware Aria OperationsVMware Cloud Foundation Operations. These symptoms are based on conditions reported for selected metrics by an external monitored system, as compared to metric symptoms, which are based on thresholds that VMware Aria OperationsVMware Cloud Foundation Operations is actively monitoring.<br>• Smart Early Warning. Add a symptom that uses a defined condition that is triggered when the number of anomalies on an object is over the trending threshold. This symptom represents the overall anomalous behavior of the object. Anomalies are based on VMware Aria OperationsVMware Cloud Foundation Operations analysis of the number of applicable metrics that violate the dynamic threshold that determines the normal operating behavior of the object. This symptom is not configurable. You either use it or you do not use it. |
| Filter by Object Type | Available only when you select a Defined On value other than Self.<br><br>Limits the symptoms to those that are configured for the selected object type based on the selected Defined On relationship. |
| Create New Symptom | If symptoms that you need for your alert do not exist, you can create them.<br><br>Opens the symptoms definition dialog box.<br><br>Not available for Smart Early Warning symptoms, which are predefined in the system. |
| All Filters | Filter the list of symptom definitions. This selection is available when Defined On is set to **Self**, or when it is set to another relationship and you select an object from the Filter by Object Type drop-down menu.<br><br>• Symptom. Type text to search on the name of the symptom definitions. For example, to display all symptom definitions that have efficiency in their name, type `Efficiency`.<br>• Defined By. Type text to search for the name of the adapter that defines the symptom definitions. For example, to display all symptom definitions provided |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | by the vCenter Adapter, type `vCenter`. To display only user-defined symptom definitions, type the search term `User`.<br>To clear a filter, click the double arrow icon that appears next to the filter name. |
| Quick filter (Name) | Search the list based on the symptom name. |
| Symptoms list | List of existing symptoms for the selected object type. To configure a symptom, drag it into the left workspace.<br><br>To combine symptoms that are based on multiple levels in the hierarchy, select the new Defined On level and Filter by Object Type before you select and drag the new symptom to the workspace. |
| Conditions tab | |
| Select Specific Object | Select a specific object based on its object type, adapter type, policy, collection state, and status. |
| Filter | Search the metrics based on object type. |
| Conditions list | List of metrics for the selected object type. To configure a condition, drag it into the left workspace. |

Use the workspace to configure the interaction of the symptoms, symptom sets, and conditions.

**Table 155: Symptom Sets in the Alert Definition Workspace**

| Option | Description |
|---|---|
| Trigger alert when {operator} of the symptom sets are true | Select the operator for all of the added symptom/condition sets. Available only when you add more than one symptom/condition set.<br><br>• All. All of the symptom/condition sets must be true before the alert is generated. Operates as a Boolean AND.<br>• Any. One or more of the symptom/condition sets must be true before the alert is generated. Operates as a Boolean OR. |
| Symptoms | The symptom/condition sets comprise an expression that is evaluated to determine if an alert should be triggered.<br><br>To add one or more symptoms from the symptom list to an existing symptom set, drag the symptom from the list to the symptom set. To create a new symptom set for the alert definition, drag a symptom to the landing area outlined with a dotted line. |
| Symptom sets | Add one or more symptoms to the workspace, define the points at which the symptom sets are true, and specify whether all or any of the symptoms in the symptom set must be true to generate the alert.<br><br>A symptom set can include one or more symptoms/conditions, and an alert definition can include one or more symptom/condition sets. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | If you create a symptom set where the Defined On object is Self, you can set the operator for multiple symptoms in the symptom set.<br><br>If you create a symptom set where the Defined On object is a relationship other than Self, you can set the operator and modify the triggering threshold. To configure the symptom set criteria, you set the options.<br>• Value operator. Specifies how the value you provide in the value text box is compared to a number of related objects to evaluate the symptom/condition set as true.<br>• Value text box. Number of objects of the specified relationship, based on the value type, that are required to evaluate the symptom/condition set as true.<br>• Value type. Possible types include the following items:<br>– Count. Exact number of related objects meet the symptom/condition set criteria.<br>– Percent. Percentage of total related objects meet the symptom/ condition set criteria.<br>– Any. One or more of the related objects meet the symptom/condition set criteria.<br>– All. All of the related objects meet the symptom/condition set criteria.<br>• Symptom set operator. Operator applied between symptoms/conditions in the symptom set.<br>– All. All of the symptoms/conditions must be true before the alert is generated. Operates as a Boolean AND.<br>– Any. One or more of the symptoms/condition must be true before the alert is generated. Operates as a Boolean OR.<br><br>When you include a symptom in a symptom set, the condition must become true to trigger the symptom set. However, you might want to configure a symptom set where the absence of a symptom condition triggers a symptom. To use the absence of the symptom condition, click the vertical ellipsis on the left of the symptom name and select **Invert Symptom**.<br><br>Although you can configure symptom criticality, if you invert a symptom, it does not have an associated criticality that affects the criticality of generated alerts. |

**Table 156: Conditions in the Alert Definition Workspace**

| Option | Description |
|---|---|
| Alert is triggered when {operator} of the sets are true | Select the operator for all of the added condition sets. Available only when you add more than one condition set.<br><br>• All. All of the condition sets must be true before the alert is generated. Operates as a Boolean AND.<br>• Any. One or more of the condition sets must be true before the alert is generated. Operates as a Boolean OR. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Conditions | The condition sets comprise an expression that is evaluated to determine if an alert should be triggered.<br>• Condition. Determines how the value you specify in the value text box is compared to the current value of the metric or property when the condition is evaluated.<br>• Value. Value that specifies the threshold.<br>• Criticality level. Severity of the symptom/condition when it is triggered.<br>• Wait Cycle. The trigger condition should remain true for this number of collection cycles before the symptom/condition is triggered. The default value is 1, which means that the symptom/condition is triggered in the same collection cycle when the condition became true.<br>**NOTE**<br>You cannot edit the wait cycle while defining conditions for Properties and Population.<br>• Cancel Cycle. The symptom/condition is canceled after the trigger condition is false for this number of collection cycles after which the symptom/condition is cancelled. The default value is 1, which means that the symptom/condition is canceled in the same cycle when the condition becomes false.<br>**NOTE**<br>You cannot edit the cancel cycle while defining conditions for Properties and Population.<br>To add one or more conditions from the condition list to an existing symptom/condition set, drag the condition from the list to the symptom/condition set. |

6. Click **Next** to add recommendations.
7. In the **Recommendations** tab, drag the selected recommendation in to the left pane. Use the workspace on the left to to change the priority order.

**Table 157: Add Recommendations Options in the Alert Definition Workspace**

| | |
|---|---|
| Create New Recommendation | If recommendations that you need to resolve the symptoms in the problem do not exist, you can create them. |
| All Filters | Filter the list of recommendations.<br>• Description. Type text to search on the name of the recommendation. For example, to display all recommendations that have memory in their name, type `Memory`.<br>• Defined By. Type text to search for the name of the adapter that defines the recommendation. For |

*Table continued on next page*

*Continued from previous page*

| | |
|---|---|
| | example, to display all recommendations provided by the vCenter Adapter, type `vCenter`.<br>To clear a filter, click the double arrow icon that appears next to the filter name. |
| Quick filter (Name) | Limits the list based on the text you enter. |
| List of available recommendations. | List of existing recommendations that you can drag to the workspace.<br><br>Recommendations are instructions and, where possible, actions that assist you with resolving alerts when they are triggered. |
| Recommendation workspace | Add one or more recommendations to the workspace.<br><br>If you add more than one recommendation, you can drag the recommendations to change the priority order. |

8. Click **Next** to activate policies.
9. In the **Policies** tab, you can view the policy tree in the left pane and you can either select the default policy or any other policy from the tree.
   You can automate the recommended action that has the highest priority by changing the **Status** to **Activated** in the right pane. Whenever the alert is executed on an object within the policy, the recommended action will be executed on the object.

   **NOTE**
   To deactivate alerts associated with a specific policy, deselect the policy in the left pane, and click **Update**.

   You can also customize thresholds for a policy by clicking the policy and editing the trigger value in the right pane. Editing the threshold of conditions will affect its alert definition in the selected policy.

   **NOTE**
   If you create an alert without enabling any policies, then the alert remains inactive.
10. Click **Create** to create the alert. The new alert appears in the list of alert definitions.

## Symptom Definitions in VMware Aria OperationsVMware Cloud Foundation Operations

Symptom Definitions
Symptoms are conditions that indicate problems in your environment. You can define symptoms in VMware Aria OperationsVMware Cloud Foundation Operations and add them to alert definitions so that you know when a problem occurs with your monitored objects.

As data is collected from your monitored objects, the data is compared to the defined symptom condition. If the condition is true, then the symptom is triggered.

To define symptoms, from the left menu, click **Operations › Configurations**, and then click the **Symptom Definitions** tile. .


You can define symptoms based on:

- Metric/Property
- Message Events
- Faults
- Metric Events

The symptoms defined in your environment are managed in the Symptom Definitions. When the symptoms that are added to an alert definition are triggered, they contribute to a generated alert.

### Define Symptoms to Cover All Possible Severities and Conditions

Use a series of symptoms to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of Warning while `Volume reached capacity limit` might have a severity level of Critical. The first symptom is not an immediate threat. The second symptom is an immediate threat.

### Understanding Negative Symptoms for VMware Aria OperationsVMware Cloud Foundation Operations Alerts

Understanding Negative Symptoms for Alerts
Understanding Negative Symptoms for Alerts
Alert symptoms are conditions that indicate problems in your environment. When you define an alert, you include symptoms that generate the alert when they become true in your environment. Negative symptoms are based on the absence of the symptom condition. If the symptom is not true, the symptom is triggered.

To use the absence of the symptom condition in an alert definition, you negate the symptom in the symptom set.

All defined symptoms have a configured criticality. However, if you negate a symptom in an alert definition, it does not have an associated criticality when the alert is generated.

All symptom definitions have a configured criticality. If the symptom is triggered because the condition is true, the symptom criticality will be the same as the configured criticality. However, if you negate a symptom in an alert definition and the negation is true, it does not have an associated criticality.

When negative symptoms are triggered and an alert is generated, the effect on the criticality of the alert depends on how the alert definition is configured.

The following table provides examples of the effect negative symptoms have on generated alerts.

**Table 158: Negative Symptoms Effect on Generated Alert Criticality**

| Alert Definition Criticality | Negative Symptom Configured Criticality | Standard Symptom Configured Criticality | Alert Criticality When Triggered |
|---|---|---|---|
| Warning | One Critical Symptom | One Immediate Symptom | Warning. The alert criticality is based on the defined alert criticality. |
| Symptom Based | One Critical Symptom | One Warning Symptom | Warning. The negative symptom has no associated criticality and the criticality of the standard symptom determines the criticality of the generated alert. |
| Symptom Based | One Critical Symptom | No standard symptom included | Info. Because an alert must have a criticality and the negative alert does not have an associated criticality, the generated alert has a criticality of Info, which is the lowest possible criticality level. |

## Recommendations in VMware Aria OperationsVMware Cloud Foundation Operations

Recommendations
Recommendations are probable solutions for an alert generated in VMware Aria OperationsVMware Cloud Foundation Operations. You can create a library of recommendations that include instructions to your environment administrators or actions that they can run to resolve an alert.

Recommendations provide your network engineers or virtual infrastructure administrators with information to resolve alerts.

Depending on the knowledge level of your users, you can provide more or less information, including the following options, in any combination.

- One line of instruction.
- Steps to resolve the alert on the target object.
- Hyperlink to a Web site, runbook, wiki, or other source.
- Action that makes a change on the target object.

When you define an alert, provide as many relevant action recommendations as possible. If more than one recommendation is available, arrange them in priority order so that the solution with the lowest effect and highest effectiveness is listed first. If no action recommendation is available, add text recommendations. Be as precise as possible when describing what the administrator should do to fix the alert.

### Where You Find Recommendations

To define recommendations, click **Operations › Configurations**, and then click the **Recommendations** tile. .

You can also define recommendations when you create an alert definition.

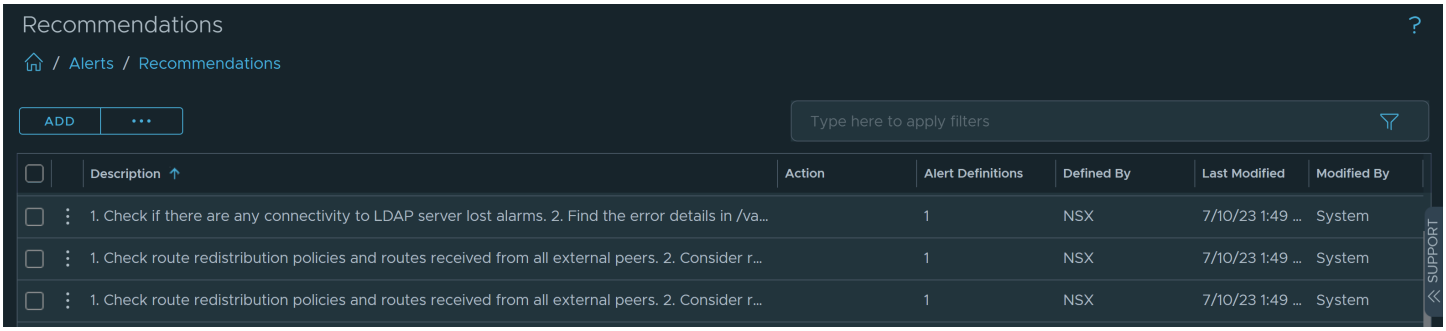


| Option | Description |
|---|---|
| Toolbar options | Use the toolbar options to manage your recommendations.<br>• Add. Add a recommendation.<br><br>Click the horizontal ellipsis to perform the following actions.<br>• Edit. Modify the selected recommendation.<br>• Delete. Remove the selected recommendation.<br>• Clone. Create a copy of the selected recommendation so that you can create a new recommendation that uses the current one.<br>• Export. Downloads the recommendations.<br>• Import. Allows you to import recommendations. To import:<br>– Click the **Import** option from the horizontal ellipsis.<br>– Click **Browse** and select the file to import.<br>– Select if you want to Overwrite or Skip the file in case of a conflict.<br>– Click **Import** to import the recommendation, and click **Done**. |
| Quick Filter | Limits the list based on the text you type.<br><br>You can also sort on the columns in the data grid. |
| Description | Displays the recommendation text that is provided when the alert is generated.<br><br>Click this link to view the Details Page. On this page, you can view the alert definitions assigned for a particular recommendation. To remove the selected recommendation from all alert definitions, click **Edit** and then, in the **Edit Recommendation** page, click **Remove from all**. |
| Action | If the recommendation includes running an action, the name of the action is displayed. |
| Alert Definitions | Displays the number of alert definitions assigned for a particular recommendation. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Defined By | Indicates whether the recommendation was created by a user or provided with a solution adapter. |
| Last Modified | Displays the date on which the recommendation was last modified. |
| Modified By | Displays the name of the user who last modified the recommendation. |

**Defining Recommendations for Alerts**

You can create recommendations that are solutions to alerts generated in VMware Aria OperationsVMware Cloud Foundation Operations. The recommendations are intended to ensure that your network operations engineers and virtual infrastructure administrators can respond to alerts as quickly and accurately as possible. A recommendation can contain links to useful Web sites or local runbooks, instructions as text, or actions that you can initiate from VMware Aria OperationsVMware Cloud Foundation Operations.

1. To define recommendations, from the left menu, click **Operations › Configurations**, and then click the **Recommendations** tile.
2. Click **Add** and enter the following details.

| Option | Description |
|---|---|
| Description | Enter the description of what must be done to resolve the triggered alert.<br><br>The description can include steps a user must take to resolve the alert or it might be instructions to notify a virtual infrastructure administrator.<br><br>This is a text field. |
| Create a hyperlink | To create a hyperlink, enter a description, select the text, and click the Hyperlink icon to make the text a hyperlink to a Website or local wiki page. |
| Action (Optional) | |
| Adapter Type | Select an adapter type from the drop-down list to narrow down the list of actions displayed in the Actions field. |
| Action | You can add an action as a method to resolve a triggered symptom or a generated alert. Actions must already be configured in VMware Aria OperationsVMware Cloud Foundation Operations.<br><br>You must provide text in the text box to describe the action before you can save the recommendation. |

3. Click **Save**.

These actions, named `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, they can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`
- `Set CPU Count for VM Power Off Allowed`
- `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

## Notifications in VMware Aria OperationsVMware Cloud Foundation Operations

Notifications
Notifications are alert notifications that meet the filter criteria in the notification rules before they are sent outside VMware Aria OperationsVMware Cloud Foundation Operations. You can configure notification rules for the supported outbound alerts so that you can filter the alerts that are sent to the selected external system.

You can use the notifications list to manage your rules and then use the notification rules to limit the alerts that are sent to the external system. To use notifications, the supported outbound alert plug-ins must be added and running.

With notification rules, you can limit the data that is sent to the following external systems.

- Standard Email. You can create multiple notification rules for various email recipients based on one or more of the filter selections. If you add recipients but do not add filter selections, all the generated alerts are sent to the recipients.
- REST. You can create a rule to limit alerts that are sent to the target REST system so that you do not need to implement filtering on that target system.
- SNMP Trap. You can configure VMware Aria OperationsVMware Cloud Foundation Operations to log alerts on an existing SNMP Trap server in your environment.
- Log File. You can configure VMware Aria OperationsVMware Cloud Foundation Operations to log alerts to a file on each of your VMware Aria OperationsVMware Cloud Foundation Operations nodes.

You configure notification options to specify which alerts are sent out for the Standard Email, REST, SNMP, and Log File outbound alert plug-ins. For the other plug-in types, all the alerts are sent when the target outbound alert plug-in is activated.

The most common outbound alert plug-in is the Standard Email plug-in. You configure the Standard Email plug-in to send notifications to one or more users when an alert is generated that meets the criteria you specify in the notification settings.

### Where You Find Notifications

To manage your notifications, from the left menu, click **Operations › Configurations**, and then click the **Notifications** tile.

> **NOTE**
> To use notifications, the supported outbound alert plug-ins must be added and running.

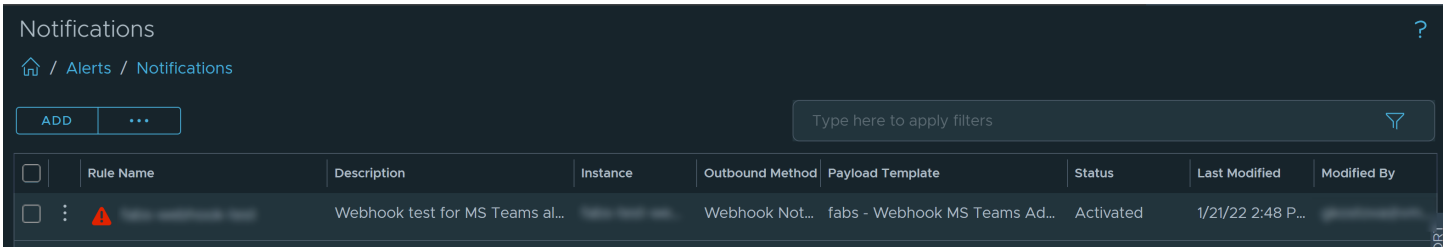


| Option | Description |
|---|---|
| Toolbar options | Use the toolbar options to manage your notification rules.<br>• Add. Opens the Add Rule dialog box where you configure the filtering options for the notification rule. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Click the horizontal ellipsis to perform the following actions.<br>• Delete. Removes the selected rule.<br>• Deactivate or Activate. Deactivates or activates the selected rule(s).<br>• Export or Import. Export the selected notifications to a ".xml" file so that you can import it on another VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Quick Filter (Action Name) | Limits the list to actions matching the filter. You can filter by:<br>• Rule Name<br>• Instance<br>• Status<br>• Modified By |
| Rule Name | Name you assigned when you created the notification rule. Click the vertical ellipsis to perform the following actions.<br>• Edit. Allows you to edit the selected rule.<br>• Clone. Allows you to clone an existing notification rule and edit the attributes of the cloned notification rule. You can create multiple alert notification rules so that you can send the same alert notifications to different outbound settings.<br>**NOTE**<br>You can clone only one alert notification rule at a time.<br>• Delete. Removes the selected rule.<br>• Deactivate or Activate. Deactivates or activates the selected rule.<br>• Export. Downloads the notification setting. |
| Description | Description of the notification rule. |
| Instance | Name of the configured outbound alert instance for the notification rule.<br>Instances are configured as part of the outbound alerts and can indicate different email servers or sender addresses for alert notifications. |
| Outbound Method | Displays the type of the outbound method that is configured. |
| Payload Template | Displays the payload template that is used. |
| Status | Displays if the rule is activated or not. |
| Email Address | If the rule is for standard email notifications, the alert recipient email addresses are listed. |
| Object Name | If the rule specifies a notification for a particular object, the object name is listed. |
| Children | If the rule specifies a notification for a particular object and selected child objects, the child object types are listed. |
| Last Modified | Displays the date on which the rule was last modified. |
| Modified By | Displays the name of the user who last modified the rule. |

### Creating Notification Rules for Alerts

You add, manage, and edit your notification rules in VMware Aria OperationsVMware Cloud Foundation Operations. To send notifications to a supported system, you must configure and activate the settings for outbound alerts.

Before you can create and manage your notification rules, you must configure the outbound alert plug-in instances. For details on configuring outbound plug-ins, see Adding Outbound Notification Plug-Ins.

You use the Notifications page to manage your alert notification rules. The rules determine which VMware Aria OperationsVMware Cloud Foundation Operations alerts are sent to the supported target systems.

Notification rules are filters that limit the data sent to external systems by using outbound alert plug-ins that are supported, configured, and running. Rather than sending all alerts to all your email recipients, you can use notification rules to send specific alerts. For example, you can send health alerts for virtual machines to one or more of your network operations engineers. You can send critical alerts for selected hosts and clusters to the virtual infrastructure administrator for those objects. Before you can create and manage notification rules, you must configure the outbound alert plug-in instances.

1. To manage your notifications, from the left menu, click **Operations › Configurations**, and then click the **Notifications** tile. On the toolbar, click **Add** to add a rule, or click the vertical ellipsis and select **Edit** to edit the selected rule.
2. Enter the following notification details.

| Option | Description |
|---|---|
| Name | Name of the rule that you use to manage the rule instance. |
| Description | Description of the rule. |
| Notification Status | Either activate or deactivate a notification setting. Deactivating a notification will stop the alert notification for that setting and activating it will activate the alert notification. |
| Advanced Settings | |
| Notification Type | Select **Alert** from the drop-down menu.<br><br>**NOTE**<br>Select **Action** as your Notification Type if you want to create Workload Placement (WLP) Action based notification. For details, see Creating Notification Rules for Notification Type 'Action'. |

3. Click **Next**.
4. Define criteria for the notification rule.

| Option | Description |
|---|---|
| Object Scope | |
| Criteria | Object Type, Object, Tags, Applications, and Tiers for which you are filtering the alert notifications.<br><br>After you select the type, you select the specific instance. For example, if you select **Object**, you then select the specific object by name and determine whether to include any child objects. |
| Alert Scope | |
| Category | Alert Types/Subtypes, Alert Impact, or Alert Definition that triggers the alert.<br><br>After you select the criteria, you can configure the specific selections associated with the criteria. For example, if you select **Alert Definition**, you then select the alert definition that limits the data to alerts with this definition. You can select multiple alert definitions as conditions for a notification to trigger. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Criticality | Defined criticality of the alert that results in the data being sent to an external system. For example, if you select **Critical**, then the data that is sent to the external system must also be labeled as critical. |
| Control State | State of the alert, either opened, assigned, or suspended. |
| Notify On | |
| Status | Current state of the alert, either canceled, updated, or new. |
| Notification Heartbeat | |
| Heartbeat | Set this to **Active** if you want to send repeat notifications for the active alerts. The frequency of the notification depends on the collection interval set for the adapter whose object is being evaluated. By default, this checkbox is not selected.<br>**NOTE**<br>Setting this option to **Active** can cause a potential surge in the notifications that are generated. |
| Advanced Filters: By Collector | |
| Collector/Group | Select a collector or group if you want to receive notifications for the objects that receive data from the selected collector/group. |
| **NOTE**<br>If you do not define any alert filters in the Define Criteria tab, then the notification will be sent for all the alerts without applying any conditions for the object scope, alert scope, or alert state. | |

5. Click **Next**.
6. Select the outbound method that you want to use to send your notification.

| Option | Description |
|---|---|
| Outbound Method | • Select Plug-In Type: Type of plugin. Select one of the outbound alert plug-in types: Log File Plugin, Rest Notification Plugin, Standard Email Plugin, SNMP Trap Plugin, Webhook Notification Plugin, Slack Plugin, and Service-Now Notification Plugin.<br>**NOTE**<br>The Rest Notification Plugin is deprecated in this release. Although you still can configure the Rest Notification Plugin, you will not be able to use a custom template for it. You can use the Webhook Notification Plugin instead of the Rest Notification Plugin. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Select Plug-In Type: Type of plugin. Select one of the outbound alert plug-in types: Standard Email Plugin, Rest Notification Plugin, Webhook Notification Plugin, Service-Now Notification Plugin, and Slack Plugin.<br>• Select Instance: Select the configured instance for the type of plug-in.<br>• Create New Instance: You can also create a new outbound instance for the plug-in type you select. For details, see Adding Outbound Notification Plug-Ins.<br><br>For details, see Adding Outbound Notification Plug-Ins. |

7. Click **Next**.
8. Select the payload template.

| Option | Description |
|---|---|
| Payload Template | Select the payload template that you want to include in the notification. Each plug-in has its default template and you can select the default template if no customization is required. The template includes additional information about the alert or the object that is displayed in the notification. You can also customize your payload for a Webhook Notification Plugin. For details on creating payload templates, see Creating Payload Templates for Outbound Plugins. |
| The values in this tab differ based on the outbound plug-in you have selected in the previous step. | |
| Outbound Method -Standard Email Plugin | If you are configuring notifications for standard email, you can add recipients and associated information.<br><br>• Recipient(s). Enter the email addresses of the individuals to whom you are sending email messages that contain alert notifications. If you are sending to more than one recipient, use a semicolon (;) between addresses.<br>• Cc Recipients. Enter the email addresses of the individuals that have to be cc'd for the email.<br>• Bcc Recipients. Enter the email addresses of the individuals that have to be bcc'd for the email.<br>• Notify again. Number of minutes between notifications messages for active alerts. Leave the text box empty to send only one message per alert.<br>• Max Notifications. Number of times to send the notification for the active alert. Leave the text box empty to send only one message per alert. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Delay to notify. Number of minutes to delay before sending a notification when a new alert is generated. For example, if the delay is 10 minutes and a new alert is generated, the notification is not sent for 10 minutes. If the alert is canceled in those 10 minutes, the notification is not sent. The notification delay reduces the number of notifications for alerts that are canceled during that time.<br>• Description. Enter the text to include in the email message. For example, the Attention Host Management team. |
| Outbound Method - Service-Now Notification Plugin | If you are configuring notifications for a Service-Now notification plug-in, you can add instances and associated information.<br>• Caller. Enter the name of the person who reported the incident or who is affected by the incident.<br>• Category. Specify the category to which the incident belongs.<br>• Sub Category. Specify the sub-category to which the incident belongs.<br>• Business Service. Specify the business service of the incident.<br>• Contact Type. Enter the contact type.<br>• State. Enter the incident state in digits.<br>• Resolution Code. Enter the resolution code for the incident.<br>• Resolution notes. Enter the resolution notes for the incident.<br>• On hold reason. Enter the reason as to why the incident is on hold.<br>• Impact. Set the incident impact in digits. Impact measures the business criticality of the affected service.<br>• Urgency. Set urgency for the incident in digits. Urgency defines the number of days taken to resolve an incident.<br>• Priority. Enter the priority for the incident. Priority defines the sequence in which the incident must be resolved.<br>• Assignment Group. Enter the assignment group for the incident.<br>• Assigned To. Enter the details of the person to whom the incident is assigned.<br>• Severity. Set the severity for the incident in digits.<br>• Upon Approval. Specify the next steps to be taken upon incident approval. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Problem. Enter the details of the related problem if it exists.<br>• Cause by change. Enter the change request which triggered the incident.<br>• Change Request. Enter the details for the related change list if it exists. |
| Outbound Method - Slack Plugin | If you are configuring notifications for a Slack plugin, add the Webhook URL of Slack. For example, the Webhook URL is in the format: `https://hooks.slack.com/services/T00000000/B00000000/XXXXXXXXXXXXXXXXXXXXXXXX`.<br>Create and authorize an app within Slack to obtain the Webhook URL. For details on creating and authorizing an app within Slack, refer to the Slack Documentation.<br><br>Once you have created the notification rule, the alerts are displayed within that particular Slack channel with a link to the alert. Click the link to view the details of the alert on the Object Summary page. |

9. Click **Next** to test the notification
10. In the **Test Notification** tab, click **Initiate Process** to initiate the notification configuration validation process.
11. If you activate the **Filter the alert definitions and objects based on criteria outlined in the Define Criteria section** option, then the alert definitions and objects displayed below are based on the criteria outlined in the **Define Criteria** section.
12. Select an alert definition and an object for validation.
13. Click **Validate Configuration**.
    – Notification Validation Steps: View the steps involved in validating the notification configuration. The steps also indicate an error, if any. The validation steps differ based on the selected Outbound method.

| Outbound Method | Validation Steps |
|---|---|
| Log File Plugin | • Validate Criteria<br>• Permissions<br>• File Created |
| Standard Email Plugin | • Validate Criteria<br>• Establish Endpoint Connection<br>• Certificates<br>• Authentication<br>• Send Notification<br><br>**NOTE**<br>The Standard Email Plugin does not provide a response that can be validated in VMware Aria OperationsVMware Cloud Foundation Operations. |
| SNMP Trap Plugin | • Validate Criteria<br>• Establish Endpoint Connection |

*Table continued on next page*

*Continued from previous page*

| Outbound Method | Validation Steps |
|---|---|
| | • Send Notification<br><br>**NOTE**<br>The SNMP Trap Plugin does not provide a response that can be validated in VMware Aria OperationsVMware Cloud Foundation Operations. |
| Webhook Notification Plugin | • Validate Criteria<br>• Establish Endpoint Connection<br>• Certificates<br>• Authentication<br>• Send Notification<br>• Endpoint Receives Notification |
| Slack Plugin | • Validate Criteria<br>• Establish Endpoint Connection<br>• Authentication<br>• Send Notification<br>• Endpoint Receives Notification |
| ServiceNow Notification Plugin | • Validate Criteria<br>• Establish Endpoint Connection<br>• Authentication<br>• Send Notification<br>• Endpoint Receives Notification |

  - Response: The Response tab in the right pane displays if the test passed successfully or if there were any errors.
  - Body: Displays the content of the notification.

14. Click **Create** to create the notification rule. You can view the rule you created under **Alerts › Notifications**.

### Creating Notification Rules for Notification Type 'Action'

You can create and manage Workload Placement (WLP) Action based notification in VMware Aria OperationsVMware Cloud Foundation Operations.

You can configure actions for which you want to receive notifications. Once configured, notifications are sent indicating the successful, failed, or timed out WLP action.

**NOTE**
The notifications are sent once the WLP Virtual Machine movement task is completed irrespective of its status.

Before you can create and manage your notification rules, you must configure the Webhook Notification Plugin. For details, see Add a Webhook Notification Plugin for Outbound Instance.

To create notification rules for WLP action:

1. From the left menu, click **Operations › Configurations**, and then click the **Notifications** tile. On the toolbar, click **Add** to add a rule, or click the vertical ellipsis and select **Clone** to clone the selected rule.

**NOTE**
You cannot change the **Notification Type** while editing a selected notification rule,

2. Enter the following notification details.

| Option | Description |
|---|---|
| Name | Name of the rule that you use to manage the rule instance. |
| Description | Description of the rule. |
| Notification Status | Either activate or deactivate a notification setting. Deactivating a notification will stop the notification for that setting and activating it will activate the notification. |
| Advanced Settings | |
| Notification Type | Select **Action** from the drop-down menu. |

3. Click **Next**.
4. Define criteria for the notification rule.

| Option | Description |
|---|---|
| Object Scope: Select set of Objects for which you want to receive notifications. | |
| Criteria | Select the Criteria as **Object** from the drop-down menu.<br><br>Search for a specific object by name and determine if you want to include any child or descendant objects, and then add one or more child/descendant objects.<br><br>The action triggers on ANY of the selected objects: |
| Notify On | |
| Status | Select the action status for which you want to receive notification. You can receive notification for **Succeeded** **Failed**, and **Timed Out** statuses. |

5. Click **Next**.
6. Select the outbound method that you want to use to send your notification.

| Option | Description |
|---|---|
| Outbound Method | • By default, the outbound method supported is **Webhook Notification Plugin**.<br>• Select Instance: Select the configured instance for the Webhook plug-in.<br>• Create New Instance: You can also create a new outbound instance for the Webhook plug-in type. For details, see Add a Webhook Notification Plugin for Outbound Instance. |

7. Click **Next**.
8. Select the payload template.

| Option | Description |
|---|---|
| Payload Template | Select the Webhook payload template that you want to include in the notification. There is a Default WLP Action Webhook Template and you can select the default template if no customization is required. You can also customize your payload for a Webhook Notification Plugin. For details on creating payload templates, see Creating Payload Templates for Outbound Plugins. |

9. Click **Next** to test the notification
10. In the **Test Notification** tab, click **Initiate Process** to initiate the notification configuration validation process.
11. If you activate the **Filter the objects based on criteria outlined in the Define Criteria section** option, then the objects displayed below are based on the criteria outlined in the **Define Criteria** section.
12. Select an object for validation.
13. Click **Validate Configuration**.
    - Notification Validation Steps: View the steps involved in validating the notification configuration. The steps also indicate an error, if any. The validation steps differ based on the selected Outbound method.

| Outbound Method | Validation Steps |
|---|---|
| Webhook Notification Plugin | • Validate Criteria<br>• Establish Endpoint Connection<br>• Certificates<br>• Authentication<br>• Send Notification<br>• Endpoint Receives Notification |

    - Response: The Response tab in the right pane displays if the test passed successfully or if there were any errors.
    - Body: Displays the content of the notification.
14. Click **Create** to create the notification rule. You can view the rule you created under **Alerts › Notifications**.

### User Scenario: Create a VMware Aria OperationsVMware Cloud Foundation Operations Email Alert Notification

User Scenario: Create an Email Alert Notification
User Scenario: Create an Email Alert Notification
As a virtual infrastructure administrator, you need VMware Aria OperationsVMware Cloud Foundation Operations to send email notifications to your advanced network engineers when critical alerts are generated for mmbhost object, the host for many virtual machines that run transactional applications, where no one has yet taken ownership of the alert.

- Ensure that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see Create an Alert Definition for Department Objects.
- Ensure that at least one instance of the Standard Email Plug-In is configured and running. See Add a Standard Email Plug-In for Outbound Alerts.

1. From the left menu, click **Operations › Configurations**, and then click the **Notifications** tile.
2. Click **Add** to add a notification rule.
3. In the **Name** text box, enter a name similar to `Unclaimed Critical Alerts for mmbhost`.

4. Set the **Notification Status**, you can either activate or deactivate a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.
5. In the **Define Criteria** tab, select the objects and alerts for which you want to receive notifications.
   a) From the **Criteria** drop-down menu, select **Object**.
   b) Locate and select the object from the list.
6. Configure the Alert Scope.
   a) From the **Category** drop-down menu, select **Alert Impact**, and from the adjacent drop-down menu, select **Health**.
   b) From the **Criticality** drop-down menu, select **Critical**.
7. In the Notify On section, select **Open** from the **Status** drop-down menu.

   The Open state indicates that no engineer or administrator has taken ownership of the alert.

8. In the **Set Outbound Method** tab, select **Standard Email Plug-In** from the **Outbound method** drop-down menu, and then select the configured instance of the email plug-in.
9. In the **Select Payload Template** tab, configure the email options.
   a) In the **Recipients** text box, enter the email addresses of the members of your advance engineering team, separating the addresses with a semi-colon (;).
   b) To send a second notification if the alert is still active after a specified amount of time, enter the number of minutes in the **Notify again** text box.
   c) Type number of notifications that are sent to users in the **Max Notifications** text box.
10. Click **Create**.

You created a notification rule that sends an email message to the members of your advance network engineering team when any critical alerts are generated for the mmbhost object and the alert is not claimed by an engineer. This email reminds them to look at the alert, take ownership of it, and work to resolve the triggering symptoms.

Respond to alert email notifications. For details, refer to the topic 'Respond to an Alert in Your Email' in the *VMware Aria Operations User Guide*.

Respond to alert email notifications.

### Notifications - User Scenario: Create a Webhook Alert Notification

As a virtual infrastructure administrator, you need VMware Aria OperationsVMware Cloud Foundation Operations to send alerts in JSON or XML to a Webhook with any endpoint REST API that accepts these messages. You want only alerts where the virtualization alerts that affect availability alert types to go to an external application. You can then use the provided information to initiate a remediation process in that application to address the problem indicated by the alert. The notification configuration limits the alerts sent to the outbound alert instance to those matching the notification criteria.

- Ensure that you have at least one alert definition for which you are sending a notification. For an example of an alert definition, see Create an Alert Definition for Department Objects.
- Ensure that at least one instance of the Webhook Notification Plugin is configured and running. See Add a Webhook Notification Plugin for Outbound Instance.

1. From the left menu, click **Operations › Configurations**, and then click the **Notifications** tile.
2. Click **Add** to add a notification rule.
3. In the **Name** text box, enter a name similar to `Virtualization Alerts for Availability`.
4. Set the **Notification Status**, you can either activate or deactivate a notification setting. Disabling a notification stops the alert notification for that setting and enabling it activates it again.
5. In the **Define Criteria** tab, select the objects and alerts for which you want to receive notifications.

a) From the **Criteria** drop-down menu, select **Object**.
b) Locate and select the object from the list.

6. Configure the Alert Scope.
   a) From the **Category** drop-down menu, select **Alert Type**, and from the **Alert Types/Subtypes** menu, select **Availability** under **Virtualization/Hypervisor Alerts**.
   b) From the **Criticality** drop-down menu, select **Warning**.
7. In the Notify On section, select **New** from the **Status** drop-down menu.

   The New status indicates that the alert is new to the system and not updated.

8. In the **Set Outbound Method** tab, select **Webhook Notification Plugin** from the **Outbound method** drop-down menu, and then select the configured instance of the Webhook plugin.
9. In the **Select Payload Template** tab, select the **Default Webhook Template**.
10. Click **Create**.

You created a notification rule that sends the alert text to the target REST-activated system. Only the alerts where the configured alert impact is Virtualization/Hypervisor Availability and where the alert is configured as a warning are sent to the target instance using the Webhook plugin.

You created a notification rule that sends the alert text to the target Webhook system.

## Outbound Settings in VMware Aria OperationsVMware Cloud Foundation Operations

Outbound Settings
You use the Outbound Settings to manage your communication settings so that you can send information to users or applications outside of VMware Aria OperationsVMware Cloud Foundation Operations.

You manage your outbound options from this page, including adding or editing outbound plug-ins, and turning the configured plug-ins on or off. When activated, the plug-in sends a message to users as email notifications, or sends a message to other applications.

Outbound plug-in settings determine how the supported external notification systems connect to their target systems. You configure one or more instances of one or more plug-in types so that you can send data about generated notifications outside of VMware Aria OperationsVMware Cloud Foundation Operations.

You configure each plug-in with the required information, including destination locations, hosts, ports, user names, passwords, instance name, or other information that is required to send notifications to those target systems. The target systems can include email recipients, log files, or other management products.

Some plug-ins are included with VMware Aria OperationsVMware Cloud Foundation Operations, and others might be added when you add a management pack as a solution.

The configuration options vary depending on which plug-in you select from the **Plug-In Type** drop-down menu.

To add outbound notification plug-in, see Adding Outbound Notification Plug-Ins.

## Where You Find Outbound Settings

To manage your outbound settings, from the left menu, click **Operations › Configurations**, and then click the **Outbound Settings** tile.

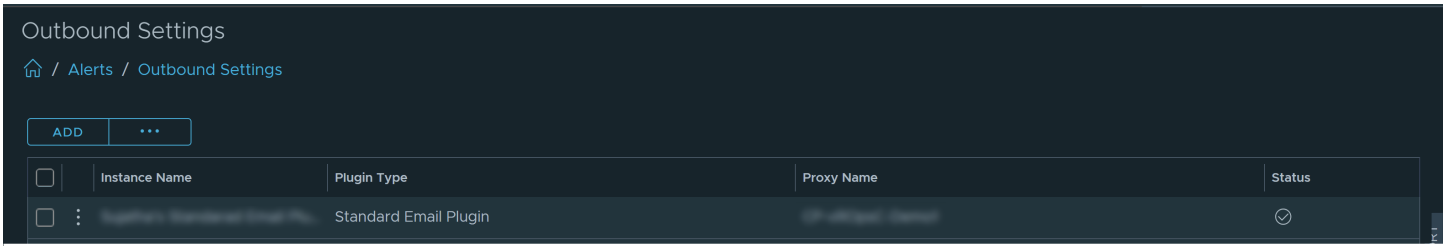


| Option | Description |
|---|---|
| Toolbar options | Use the toolbar options to manage your Outbound Plug-Ins.<br>• Add. Opens the Outbound Plug-In dialog box where you configure the connection options for the instance.<br><br>Select an existing plugin and click the vertical ellipsis to perform the following actions.<br><br>• Edit. Modify the Outbound Plug-In instance details.<br>• Delete. Removes the selected plug-in instance.<br>• Activate or Deactivate. Starts or stops the plug-in instance. Deactivating an instance allows you to stop sending the messages configured for the plug-in without removing the configuration from your environment.<br>• Export. Downloads the outbound settings.<br>• Import. Allows you to import outbound settings. To import:<br>– Click the **Import** option from the horizontal ellipsis.<br>– Click **Browse** and select the file to import.<br>– Select if you want to Overwrite or Skip the file in case of a conflict.<br>– Click **Import** to import outbound settings, and click **Done**. |
| Instance Name | Name that you assigned when you created the plug-in instance.<br><br>Click the vertical ellipsis to perform the following actions.<br>• Edit. Allows you to edit the selected payload template.<br>• Delete. Removes the selected payload template.<br>• Activate or Deactivate. Starts or stops the plug-in instance. Deactivating an instance allows you to stop sending the messages configured for the plug-in without removing the configuration from your environment.<br>• Export. Downloads the outbound settings. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Plug-In Type | Type of configured plug-in for the plug-in instance. The types of plug-ins vary depending on the solutions you have added to your environment.<br><br>The most common plug-in types include standard email, SNMP trap, log file, and REST. |
| Status | Specifies whether the plug-in is currently running. |

**List of Outbound Plugins**

VMware Aria OperationsVMware Cloud Foundation Operations provides outbound plug-ins. This list includes the name of the plug-in and whether you can filter the outbound data based on your notification settings.

If the plug-in supports configuring notification rules, then you can filter the messages before they are sent to the target system. If the plug-in does not support notifications, all messages are sent to the target system, and you can process them in that application.

If you installed other solutions that include other plug-in options, they appear as a plug-in option with the other plug-ins.

Messages and alerts are sent only when the plug-in is activated.

**Table 159: Notification Support for Outbound plug-ins**

| Outbound plug-in | Configure Notification Rules |
|---|---|
| Automated Action plug-in | No<br>The Automated Action plug-in is activated by default. If automated actions stop working, select the Automated Action plug-in and activate it if necessary. If you edit the Automated Action plug-in, you only have to provide the instance name. |
| Log File plug-in | Yes<br>To filter the log file alerts, you can either configure the file named `TextFilter.xml` or configure the notification rules. |
| REST Notification plug-in | Yes |
| Network Share plug-in | No |
| Standard Email plug-in | Yes |
| SNMP Trap plug-in | Yes |
| Webhook Notification Plugin | Yes |
| Slack plug-in | Yes |
| Service-Now Notification plug-in | Yes |
| Vmware Hosted Email plug-in | No<br>The Vmware Hosted Email plug-in is activated by default. |

### Configuring HTTP Proxy for Outbound Settings

You can configure an HTTP proxy for outbound communication over HTTP/HTTPS protocol in VMware Aria Operations. Once the HTTP proxy is configured for an outbound setting, all the corresponding outbound HTTP(S) communication must happen through that proxy.

1. From the left menu, click **Operations › Configurations**, and then click the **Outbound Settings** tile. Click the **HTTP Proxy for Outbound Settings** tab.

2. Click **Add**.

**Table 160: HTTP Proxy Options**

| Options | Description |
|---|---|
| Proxy Name | Name of the HTTP proxy server. |
| Proxy Host/ IP | The IP address of the HTTP proxy. |
| Proxy Port | Port number used to connect to the HTTP proxy server. |
| Proxy Username | Username of the HTTP proxy server. |
| Proxy Password | Password for the HTTP proxy server username. |

3. Click **Save**.
   The HTTP proxy setting is added.

Use the HTTP Proxy to configure outbound plugins for Service-Now Notification Plugin, Webhook Notification Plugin, and Slack Plugin in VMware Aria Operations. For more information see Outbound Settings in .

### Importing and Exporting HTTP Proxy for Outbound Settings

You can import or export the HTTP Proxy settings for outbound plugins.

#### Importing HTTP Proxy for Outbound Settings

1. From the **HTTP Proxy for Outbound Settings** tab, click the horizontal ellipsis, and then click **Import** to import the outbound proxy settings.
2. Click **Browse** to select the outbound proxy setting and enter the **Encryption Key**.
3. In case of a conflict, click **Overwrite HTTP Proxy for Outbound Settings** to delete the existing proxy setting and proceed. Optionally, click **Skip HTTP Proxy for Outbound Settings** to cancel the import.

   **NOTE**
   A conflict happens when you try to import a proxy with a name that already exists in the proxy list. Proxy names are unique and cannot be repeated.

4. Click **Import**.

The HTTP proxy import process begins.

#### Exporting HTTP Proxy for Outbound Settings

1. From the **HTTP Proxy for Outbound Settings** tab, select the HTTP Proxy you want to export, click the horizontal ellipsis, and then click **Export**.
2. Enter a new password in the **Setup a new password to export data** field.
3. Re-enter the password in the **Repeat a password** field.
4. Click **Export**.

The outbound proxy settings data gets exported in the `.json` format. The passphrase entered at the time of export is used to encrypt the sensitive information. The same passphrase must be used at the time of import.

### Exporting and Importing Outbound Settings

As a VMware Cloud Foundation Operations admin, you can backup the content before upgrading, export all the outbound plugin configurations, and import it into a different VMware Cloud Foundation Operations instance. You can also export the content from VMware Cloud Foundation Operations on-prem to VMware Cloud Foundation Operations.

**NOTE**
Any user with "Manage" Outbound Settings permission can export and import outbound plugin configurations.

1. Export an outbound setting.
   a) From the left menu, click **Operations › Configurations**, and then click the **Outbound Settings** tile.
   b) Select the outbound settings that you want to export and click the horizontal ellipses and select **Export**.
   c) Setup a new password to export data. The password should be at least 14 characters long.
   d) Click **Export**.
   The outbound setting data is exported in the .json format. A password is used to encrypt the data in the file using the AES algorithm with 128 bit key. Use the same password while importing this file.
2. Import an outbound setting.

   **NOTE**
   Before importing the outbound setting, ensure that you have exported the outbound plugin configurations.

   a) From the left menu, click **Operations › Configurations**, and then click the **Outbound Settings** tile.
   b) Click the horizontal ellipses and select **Import**.
   c) Click **Browse** to select the `.json` file and enter the password that you had set while exporting the content.
   d) If there is a conflict while importing the content, you can either overwrite the existing outbound settings or skip the import, which is the default.
   e) Click **Import** to import outbound settings to the destination setup.

      **NOTE**
      While importing outbound settings on VMware Cloud Foundation Operations, HTTP Proxy configurations will be excluded.

      **NOTE**
      While importing outbound settings on VMware Cloud Foundation Operations, Cloud Proxy configurations will be excluded.

## Payload Templates in VMware Aria OperationsVMware Cloud Foundation Operations

Payload Templates
Payload is an essential information in the data block that you send or receive from the server. Use the **Payload Templates** page to view the list of payload templates available for each plug-in in VMware Aria OperationsVMware Cloud Foundation Operations. You can add, manage, and edit your payload templates from this page. Default payload templates are provided for each plug-in type.

## Where You Find Payload Templates

To manage your payload templates, from the left menu, click **Operations › Configurations**, and then click the **Payload Templates** tile.

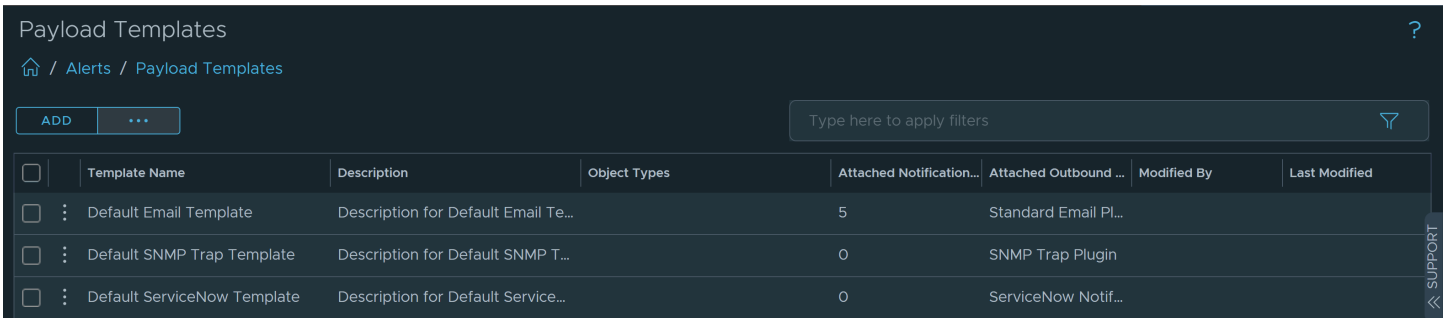


| Option | Description |
|---|---|
| Toolbar options | Use the toolbar options to manage your notification rules.<br>• Add. Use the Create Payload Template dialog box to create new payload templates.<br>Click the horizontal ellipsis to perform the following actions.<br>• Delete. Removes the selected payload template.<br>• Export. Downloads the payload template.<br>**NOTE**<br>Export and delete actions are not supported for the default payload templates available for each plug-in.<br>• Import. Allows you to import payload templates. To import:<br>– Click the **Import** option from the horizontal ellipsis.<br>– Click **Browse** and select the file to import.<br>– Select if you want to Overwrite or Skip the import in case of a conflict.<br>– Click **Import** to import the payload template, and then click **Done**. A message with the number of imported and skipped files will appear. |
| Quick Filter | Limits the list based on the text you type. It considers only text from templates Name column. You can also sort columns in the data grid. You can view a blue arrow next to the column according to which sorting was performed, pointing up or down based on the sorting order (ascending or descending). |
| Template Name | Name of the payload template.<br>Click the vertical ellipsis to perform the following actions.<br>• Edit. Allows you to edit the selected payload template.<br>• Clone. Clones the selected payload template.<br>• Delete. Removes the selected payload template.<br>• Export. Downloads the payload template.<br>•<br>**NOTE**<br>Edit, Delete, and Export actions are not supported for the default payload templates available for each plug-in. |
| Description | Description of the payload template. |
| Object Types | Base object type against which the payload template is defined, if any. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Attached Notification Rules | Notification rule attached to the payload template. |
| Attached Outbound Methods | Outbound plugin type attached to the payload template. |
| Modified By | Name of the last person to modify the payload template. |
| Last Modified | Date on which the payload template was last modified. |

**Creating Payload Templates for Outbound Plugins**

You can create a payload template for an outbound plugin of your choice in VMware Aria OperationsVMware Cloud Foundation Operations.

Use payload templates to configure the payload of an email, and customize the subject line, and email body. You can enable your own input properties and different payloads for updated and canceled alerts.
You can customize payloads only for Standard Email Plug-in and Webhook Notification Plug-in.

1. From the left menu, click **Operations › Configurations**, and then click the **Payload Templates** tile. On the toolbar, click **Add** to create a new payload template.
2. In the **Details** tab, enter the basic details of the payload template.

| Option | Description |
|---|---|
| Name | Provide a name for the payload template. |
| Description | Enter a description for the payload template. |
| Outbound Method | Outbound plugin for which you want to create a new payload template.<br>Select one of the outbound alert plug-in types: Log File Plugin, Standard Email Plugin, SNMP Trap Plugin, Webhook Notification Plugin, Slack Plugin, and Service-Now Notification Plugin.<br><br>Select Plug-In Type: Type of plugin. Select one of the outbound alert plug-in types: Standard Email Plugin, Webhook Notification Plugin, Service-Now Notification Plugin, and Slack Plugin. |
| Advanced Settings | |
| Notification Type | • Select **Alert** to configure an alert notification.<br>• Select **Action** to configure notification for WLP VM movement task.<br><br>**NOTE**<br>The **Action** option appears only when you select **Webhook Notification Plugin** in the **Outbound Method** field. You can configure notification for type 'Action' only for the Webhook Notification Plugin. |

3. Click **Next**.
4. In the **Object Content** tab, define the object details that you want to include in the notifications.

**NOTE**
To create a payload template, it is mandatory to add the object type for all the outbound plugin types except for the Standard Email Plugin and Webhook Notification Plugin.

| Option | Description |
|---|---|
| Add Object Type | • Select an object type from the list. Once you select the object type, define self metrics/properties, ancestors, descendants, ancestor metrics/properties, and descendant metrics/properties associated with the object type that you want to include in the notification.<br>From the right pane, double-click or drag the metrics and properties into the **Add Metrics and Properties** box. You can select up to 30 metrics and properties.<br>**NOTE**<br>For Action Type of Notifications, only Virtual Machine is available for selection.<br>• Define the ancestor/descendant object type along with the corresponding metrics and properties. Select the **ancestor/descendant** from the drop-down menu and from the right pane, double-click or drag the corresponding metrics and properties into the left pane.<br>**NOTE**<br>For SNMP Trap Plugin, you can only select ancestor objects but you cannot add metrics/properties or define descendant details.<br>The information that you define here will be included in the alert notification for all the plug-ins. However, for a Webhook Notification Plugin and Standard Email Plugin, the information will be included only when you define the values in the **Payload Details** tab. |

5. Click **Create** to create the new payload template or click **Next** if you are creating a payload template for a Standard Email Plugin or a Webhook Notification Plugin.
6. In the **Payload Details** tab, enter the payload details that you want to include in the notification.

**NOTE**
This tab is available only when you are creating a payload template for a Standard Email Plugin or a Webhook Notification Plugin.

| Option | Description |
|---|---|
| Do you want to add template input properties? | Select **Yes** to add input properties and enter the **Key**, **Type**, **Display Name**, and **Description** of the input property. Otherwise, select **No**.<br><br>**NOTE**<br>The input properties are specific to your endpoint. Once you define the input properties in the template, you must provide the appropriate values in each rule where this template will be used. |
| Do you want different payload details for new, updated, and canceled alerts? | Select **Yes** to define different payload details for new, updated, and canceled alerts. Otherwise, select **No**.<br><br>**NOTE**<br>This field does not appear when the **Notification Type** is **Action**. |
| The following fields appear while creating a payload template for the Standard Email Plugin. | |
| Subject | Enter a subject for the email notification. |
| Body | Enter the content for the email notification. You can also search for parameters in the right pane. Click the copy icon next to the parameter to copy the parameter and you can paste the parameter in the email body.<br>You can use the options in the toolbar to edit, format, and highlight the email content.<br><br>**NOTE**<br>You can set up different email content for new, updated, and canceled alerts. |
| The following fields appear while creating a payload template for the Webhook Notification Plugin. | |
| Endpoint URL | Enter the endpoint URL. The endpoint URL will be appended to the base URL provided in the related webhook outbound instance.<br><br>**NOTE**<br>The entire URL is encoded. However, there is an exception to use the character '/' in the URL. |
| Content Type | Select the content type for the payload. |
| Custom Headers | Enter the **HTTP Custom Header Name** and **Value**. Click the plus icon to add multiple custom headers.<br><br>**NOTE**<br>For webhook payloads using token-based authentication, add an Authorization Header in the format required by the endpoint. |
| HTTP Method | Select the HTTP method of request. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Payload of the request | Payload for the selected plug-in type. It displays information based on the selected metrics, properties, ancestors, and object types.<br>You can search for parameters in the right pane. Click the copy icon next to the parameter to copy the parameter and you can paste the parameter in the **Payload of the request** box. |

7. Click **Create**.
   Once the payload template is created, you can view it in the **Payload Templates** page. After selecting a payload template in the notification rule, you can view the payload template details in the **Notifications** page.

## Managing Alert Groups in VMware Aria OperationsVMware Cloud Foundation Operations

Managing Alert Groups
For easy and better management of alerts, you can arrange them as a group as per your requirement.

It is complicated to identify a problem in large environments as you receive different kind of alerts. To manage alerts easily, group them by their definitions.

For example, there are 1000 alerts in your system. To identify different types of alerts, group them based on their alert definitions. It is also easy to detect the alert having the highest severity in the group.

When you group alerts, you can view the number of times the alerts with the same alert definition are triggered. By grouping alerts, you can perform the following tasks easily and quickly:

- Find the noisiest alert: The alert that has triggered maximum number of times is known as the noisiest alert. Once you find it, you can deactivate it to avoid further noise.
- Filter alerts: You can filter alerts based on a substring in alert definitions. The result shows the group of alerts that contain the substring.

**NOTE**

- If you cancel or deactivate an alert group, the alerts are not canceled instantly. It might take some time if the group is large.
- Only one group can be expanded at a time.
- The number next to the group denotes the number of alerts in that particular group.
- The criticality sign  indicates the highest level of severity of an alert in a group.

### Grouping Alerts in VMware Aria OperationsVMware Cloud Foundation Operations

You can group alerts by time, criticality, definition, and object type.

To group alerts:

1. From the left menu, click **Operations › Alerts**.
2. Select from the various options available from the **Group By** drop-down menu.

### Deactivating Alerts in VMware Aria OperationsVMware Cloud Foundation Operations

In an alerts group, you can deactivate an alert by a single click.

To deactivate an alert:

1. From the left menu, click **Operations › Alerts**.
2. From the **Group By** drop-down, select **Definition**, and click on the name of the Alert Definition Group.
3. From the data grid, click **Actions › Deactivate**.
   You can deactivate the alerts by two methods:
   - Deactivate Alert in All Policies: Deactivates the alert for all the objects for all the policies.
   - Deactivate Alert in Selected Policies: Deactivates the alert for the objects having the selected policy.

## Alert Definition Best Practices

As you create alert definitions for your environment, apply consistent best practices so that you optimize alert behavior for your monitored objects.

### Alert Definitions Naming and Description

The alert definition name is the short name that appears in the following places:

- In data grids when alerts are generated
- In outbound alert notifications, including the email notifications that are sent when outbound alerts and notifications are configured in your environment

Ensure that you provide an informative name that clearly states the reported problem. Your users can evaluate alerts based on the alert definition name.

The alert definition description is the text that appears in the alert definition details and the outbound alerts. Ensure that you provide a useful description that helps your users understand the problem that generated the alert.

### Wait and Cancel Cycle

The wait cycle setting helps you adjust for sensitivity in your environment. The wait cycle for the alert definition goes into effect after the wait cycle for the symptom definition results in a triggered symptom. In most alert definitions you configure the sensitivity at the symptom level and configure the wait cycle of alert definition to 1. This configuration ensures that the alert is immediately generated after all of the symptoms are triggered at the desired symptom sensitivity level.

The cancel cycle setting helps you adjust for sensitivity in your environment. The cancel cycle for the alert definition goes into affect after the cancel cycle for the symptom definition results in a cancelled symptom. In most definitions you configure the sensitivity at the symptom level and configure the cancel cycle of alert definition to 1. This configuration ensures that the alert is immediately cancelled after all of the symptoms conditions disappear after the desired symptom cancel cycle.

### Create Alert Definitions to Generate the Fewest Alerts

You can control the size of your alert list and make it easier to manage. When an alert is about a general problem that can be triggered on a large number of objects, configure its definition so that the alert is generated on a higher level object in the hierarchy rather than on individual objects.

As you add symptoms to your alert definition, do not overcrowd a single alert definition with secondary symptoms. Keep the combination of symptoms as simple and straightforward as possible.

You can also use a series of symptom definitions to describe incremental levels of concern. For example, `Volume nearing capacity limit` might have a severity value of Warning while `Volume reached capacity limit` might have a severity level of Critical. The first symptom is not an immediate threat, but the second one is an immediate threat. You can then include the Warning and Critical symptom definitions in a single alert definition with an Any condition and set the alert criticality to be Symptom Based. These settings cause the alert to be generated with the right criticality if either of the symptoms is triggered.

### Avoid Overlapping and Gaps Between Alerts

Overlaps result in two or more alerts being generated for the same underlying condition. Gaps occur when an unresolved alert with lower severity is canceled, but a related alert with a higher severity cannot be triggered.

A gap occurs in a situation where the value is <=50% in one alert definition and >=75% in a second alert definition. The gap occurs because when the percentage of volumes with high use falls between 50 percent and 75 percent, the first problem cancels but the second does not generate an alert. This situation is problematic because no alert definitions are active to cover the gap.

**Actionable Recommendations**

If you provide text instructions to your users that help them resolve a problem identified by an alert definition, precisely describe how the engineer or administrator should fix the problem to resolve the alert.

To support the instructions, add a link to a wiki, runbook, or other sources of information, and add actions that you run from VMware Aria OperationsVMware Cloud Foundation Operations on the target systems.

## Create a Simple Alert Definition

While troubleshooting, you can now quickly create an alert for a particular object type or a metric in a quick and efficient way.

You can create a simple alert definition from the following locations.

- From the left menu, click **Operations ›  Troubleshoot** and select the metric for which you want to create an alert. You can create an alert from the **Potential Evidence** or the **Metrics** tab.
- From the left menu, click **Operations › Alerts**. Select an alert and click the **Potential Evidence** tab.
- From the Home page, you can search for the specific object type or metric in the search bar.

1. Click the drop-down menu available in the right side of the widget and select the **Create an Alert Definition** option.

   **NOTE**
   You cannot create alert definitions for Badge, Time, and Capacity Remaining metrics.

2. In the Create Alert Definition page, enter the **Name** and **Description** of the alert.
3. Set thresholds, criticality, and the number of wait cycles. Click **Show Advanced Settings** to set Wait Cycle and Cancel Cycle.

   **NOTE**
   The Object Type or Metric/Property are pre-selected and cannot be edited.

4. Click **Create**.
   The new alert is created and the policy the object belongs to and its children policies are activated for the alert.

## Create a New Alert Definition

Based on the root cause of the problem, and the solutions that you used to fix the problem, you can create a new alert definition for VMware Aria OperationsVMware Cloud Foundation Operations to alert you. When the alert is triggered on your host system, VMware Aria OperationsVMware Cloud Foundation Operations alerts you and provides recommendations on how to solve the problem.

To alert you before your host systems experience critical capacity problems, and have VMware Aria OperationsVMware Cloud Foundation Operations notify you of problems in advance, you create alert definitions, and add symptom definitions to the alert definition.

1. From the left menu,click **Operations › Configurations**, and then click the **Alert Definitions** tile.
2. Enter `capacity` in the search text box.
   Review the available list of capacity alert definitions. If a capacity alert definition does not exist for host systems, you can create one.
3. Click **Add** to create a new capacity alert definition for your host systems.

a) In the alert definition workspace, for the Name and Description, enter `Hosts - Alert on Capacity Exceeded`.
b) For the Base Object Type, select **vCenter Adapter > Host System**
c) Under **Advanced Settings**, select the following options.

| Option | Selection |
|---|---|
| Impact | Select **Risk**. |
| Criticality | Select **Immediate**. |
| Alert Type and Subtype | Select **Application : Capacity**. |
| Wait Cycle | Select **1**. |
| Cancel Cycle | Select **1**. |

d) In the **Symptoms/Conditions** workspace, select the following options.

| Option | Selection |
|---|---|
| Defined On | Select **Self**. |
| Symptom Definition Type | Select **Metric / Property**. |
| Quick filter (Name) | Enter `capacity`. |

e) From the Symptom Definition list, click **Host System Capacity Remaining is moderately low** and drag it to the left pane.
In the Symptoms pane, make sure that the Base object exhibits criteria is set to **All** by default.
f) For Add Recommendations, enter `virtual machine` in the quick filter text box.
g) Click **Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the system**, and drag it to the recommendations area in the left pane.
This recommendation is set to Priority 1.

4. Click **Save** to save the alert definition.
Your new alert appears in the list of alert definitions.

You have added an alert definition to have VMware Aria OperationsVMware Cloud Foundation Operations alert you when the capacity of your host systems begins to run out.

## Create an Alert Definition for Department Objects

As a virtual infrastructure administrator, you are responsible for the virtual machines and hosts that the accounting department uses. You can create alerts to manage the accounting department objects.

You received several complaints from your users about delays when they are using their accounting applications. Using VMware Aria OperationsVMware Cloud Foundation Operations, you identified the problem as related to CPU allocations and workloads. To better manage the problem, you create an alert definition with tighter symptom parameters so that you can track the alerts and identify problems before your users encounter further problems.

Using this scenario, you create a monitoring system that monitors your accounting objects and provides timely notifications when problems occur.

**Add Description and Base Object to Alert Definition**

To create an alert to monitor the CPUs for the accounting department virtual machines and monitor host memory for the hosts on which they operate, you begin by describing the alert.

When you name the alert definition and define alert impact information, you specify how the information about the alert appears in VMware Aria OperationsVMware Cloud Foundation Operations. The base object is the object around which the alert definition is created. The symptoms can be for the base object and for related objects.

1. From the left menu, click **Operations › Configurations**, and then click the **Alert Definitions** tile.
2. Click **Add** to add a definition.
3. Type a name and description.

   In this scenario, type `Acct VM CPU early warning` as the alert name, which is a quick overview of the problem. The description, which is a detailed overview. should provide information that is as useful as possible. When the alert is generated, this name and description appears in the alert list and in the notification.

4. From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Host System**.

   This alert is based on host systems because you want an alert that acts as an early warning to possible CPU stress on the virtual machines used in the accounting department. By using host systems as the based object type, you can respond to the alert symptom for the virtual machines with bulk actions rather than responding to an alert for each virtual machine.

5. Click **Advanced Settings** and configure the metadata for this alert definition.
   a) From the **Impact** drop-down menu, select **Risk**.

      This alert indicates a potential problem and requires attention in the near future.

   b) From the **Criticality** drop-down menu, select **Immediate**.

      As a Risk alert, which is indicative of a future problem, you still want to give it a high criticality so that it is ranked for correct processing. Because it is designed as an early warning, this configuration provides a built-in buffer that makes it an immediate risk rather than a critical risk.

   c) From the **Alert Type and Subtype** drop-down menu, select **Performance** under **Virtualization/ Hypervisor**.
   d) To ensure that the alert is generated during the first collection cycle after the symptoms become true, set the **Wait Cycle** to `1`.
   e) To ensure that the an alert is removed as soon as the symptoms are no longer triggered, set the **Cancel Cycle** to `1`.

      The alert is canceled in the next collection cycle if the symptoms are no long true.

   These alert impact options help you identify and prioritize alerts as they are generated.

You started an alert definition where you provided the name and description, selected host system as the base object type, and defined the data that appears when the alert generated.

Continue in the workspace, adding symptoms to your alert definition. See Add a Virtual Machine CPU Usage Symptom to the Alert Definition.

### Add a Virtual Machine CPU Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add symptoms to your VMware Aria OperationsVMware Cloud Foundation Operations alert definition after you provide the basic descriptive information for the alert. The first symptom you add is related to CPU usage on virtual machines. You later use a policy and group to apply alert to the accounting virtual machines.

Begin configuring the alert definition. See Add Description and Base Object to Alert Definition.

This scenario has two symptoms, one for the accounting virtual machines and one to monitor the hosts on which the virtual machines operate.

1. In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Next** and configure the symptoms.
2. Begin configuring the symptom set related to virtual machines CPU usage.

a) From the **Select Symptom** drop-down menu, select **Metric / Property**.
b) From the **Defined On** drop-down menu, select **Child**.
c) From the **Filter by Object Type** drop-down menu, select **Virtual Machine**.
d) Click **Create New** to open the **Add Symptom Definition** workspace window.

3. Configure the virtual machine CPU usage symptom in the **Add Symptom Definition** workspace window.
   a) From the **Base Object Type** drop-down menu, expand **vCenter Adapter** and select **Virtual Machine**.
   The collected metrics for virtual machines appears in the list.
   b) In the metrics list **Search** text box, which searches the metric names, type `usage`.
   c) In the list, expand **CPU** and drag **Usage (%)** to the workspace on the left.
   d) From the threshold drop-down menu, select **Dynamic Threshold**.

      Dynamic thresholds use VMware Aria OperationsVMware Cloud Foundation Operations analytics to identify the trend metric values for objects.

   e) In the **Symptom Definition Name** text box, type a name similar to `VM CPU Usage above trend`.
   f) From the criticality drop-down menu, select **Warning**.
   g) From the threshold drop-down menu, select **Above Threshold**.
   h) Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

      This Wait Cycle setting requires the symptom condition to be true for 3 collection cycles before the symptom is triggered. This wait avoids triggering the symptom when there is a short spike in CPU usage.

   i) Click **Save**.

   The dynamic symptom, which identifies when the usage is above the tracked trend, is added to the symptom list.
4. In the **Alert Definition Workspace** window, drag **VM CPU Usage above trend** from the symptom definition list to the symptom workspace on the left.
   The Child-Virtual Machine symptom set is added to the symptom workspace.
5. In the symptoms set, configure the triggering condition so that when the symptom is true on half of the virtual machines in the group to which this alert definition is applied, the symptom set is true.
   a) From the value operator drop-down menu, select **>**.
   b) In the value text box, enter `50`.
   c) From the value type drop-down menu, select **Percent**.

You defined the first symptom set for the alert definition.

Add the host memory usage symptom to the alert definition. See Add a Host Memory Usage Symptom to the Alert Definition.

### Add a Host Memory Usage Symptom to the Alert Definition

To generate alerts related to CPU usage on your accounting virtual machines, you add a second symptom to your VMware Aria OperationsVMware Cloud Foundation Operations alert definition after you add the first symptom. The second symptom is related to host memory usage for the hosts on which the accounting virtual machines operate.

Add the virtual machine CPU usage symptom. See Add a Virtual Machine CPU Usage Symptom to the Alert Definition.

1. In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, and **Alert Impact**, click **Next**.
2. Configure the symptom related to host systems for the virtual machines.
   a) From the **Select Symptom** drop-down menu, select **Metric / Property**.
   b) From the **Defined On** drop-down menu, select **Self**.
   c) Click **Create New** to add new symptom.

3. Configure the host system symptom in the **Add Symptom Definition** workspace window.
   a) From the **Base Object Type** drop-down menu, expand **vCenter Adapters** and select **Host System**.
   b) In the metrics list, expand **Memory** and drag **Usage (%)** to the workspace on the left.
   c) From the threshold drop-down menu, select **Dynamic Threshold**.

      Dynamic thresholds use VMware Aria OperationsVMware Cloud Foundation Operations analytics to identify the trend metric values for objects.

   d) In the **Symptom Definition Name** text box, enter a name similar to `Host memory usage above trend`.
   e) From the criticality drop-down menu, select **Warning**.
   f) From the threshold drop-down menu, select **Above Threshold**.
   g) Leave the **Wait Cycle** and **Cancel Cycle** at the default values of 3.

      This Wait Cycle setting requires the symptom condition to be true for three collection cycles before the symptom is triggered. This wait avoids triggering the symptom when a short spike occurs in host memory usage.

   h) Click **Save**.

   The dynamic symptom identifies when the hosts on which the accounting virtual machines run are operating above the tracked trend for memory usage.

   The dynamic symptom is added to the symptom list.
4. In the **Alert Definition Workspace** window, drag **Host memory usage above trend** from the symptoms list to the symptom workspace on the left.
   The Self-Host System symptom set is added to the symptom workspace.
5. On the Self-Host System symptom set, from the value type drop-down menu for **This Symptom set is true when**, select **Any**.

   With this configuration, when any of the hosts running accounting virtual machines exhibit memory usage that is above the analyzed trend, the symptom condition is true.
6. At the top of the symptom set list, from the **Match {operator} of the following symptoms** drop-down menu, select **Any**.
   With this configuration, if either of the two symptom sets, virtual machine CPU usage or the host memory, are triggered, an alert is generated for the host.

You defined the second symptom set for the alert definition and configured how the two symptom sets are evaluated to determine when the alert is generated.

Add recommendations to your alert definition so that you and your engineers know how to resolve the alert when it is generated. See Add Recommendations to the Alert Definition.

### Add Recommendations to the Alert Definition

To resolve a generated alert for the accounting department's virtual machines, you provide recommendations so that you or other engineers have the information you need to resolve the alert before your users encounter performance problems.

Add symptoms to your alert definition. See Add a Host Memory Usage Symptom to the Alert Definition.

As part of the alert definition, you add recommendations that include actions that you run from VMware Aria OperationsVMware Cloud Foundation Operations and instructions for making changes in vCenter that resolve the generated alert.

1. In the **Alert Definition Workspace** window, after you configure the **Name and Description**, **Base Object Type**, **Alert Impact**, and **Add Symptom Definitions**, click **Next** and add the recommended actions and instructions.

2. Click **Create New Recommendation** and select an action recommendation to resolve the virtual machine alerts.
   a) In the **Description** text box, enter a description of the action similar to `Add CPUs to virtual machines`.
   b) From the **Actions** drop-down menu, select **Set CPU Count for VM**.
   c) Click **Create**.
3. Click **Create New Recommendation** and provide an instructive recommendation to resolve host memory problems similar to this example.
   `If this host is part of a DRS cluster, check the DRS settings to verify that the load balancing setting are configured correctly. If necessary, manually vMotion the virtual machines.`
4. Click **Create**.
5. Click **Create New Recommendation** and provide an instructive recommendation to resolve host memory alerts.
   a) Enter a description of the recommendation similar to this example.
      `If this is a standalone host, add more memory to the host.`
   b) To make the URL a hyperlink in the instructions, copy the URL, for example, https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html, to your clipboard.
   c) Highlight the text in the text box and click the hyperlink icon.
   d) Paste the URL in the **Create a hyperlink** text box and click **OK**.
   e) Click **Create**.
6. In the **Alert Recommendation Workspace**, drag **Add CPUs to virtual machines**, **If this host is part of a DRS cluster**, and the **If this is a standalone host** recommendations from the list to the recommendation workspace in the order presented.
7. Click **Next** to select policies and view notifications.
8. Click **Create**.

You provided the recommended actions and instructions to resolve the alert when it is generated. One of the recommendations resolves the virtual machine CPU usage problem and the other resolves the host memory problem.

Create a group of objects to use to manage your accounting objects. See Create a Custom Accounting Department Group.

### Create a Custom Accounting Department Group

To manage, monitor, and apply policies to the accounting objects as a group, you create a custom object group.

Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.

1. From the left menu, click **Operations › Configurations**, and then click the **Custom Groups** tile under Logical Groupings.
2. Click **Add** to create a new custom group.
3. Type a name similar to `Accounting VMs and Hosts`.
4. From the **Group Type** drop-down menu, select **Department**.
5. From the **Policy** drop-down menu, select **Default Policy**.

   When you create a policy, you apply the new policy to the accounting group.

6. In the Define membership criteria area, from the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Host System**, and configure the dynamic group criteria.
   a) From the criteria drop-down menu, select **Relationship**.
   b) From the relationships options drop-down menu, select **Parent of**.
   c) From the operator drop-down menu, select **contains**.

   d) In the **Object name** text box, enter `acct`.
   e) From the navigation tree drop-down list, select **vSphere Hosts and Clusters**.

   You created a dynamic group where host objects that are the host for virtual machines with acct in the virtual machine name are included in the group. If a virtual machine with acct in the object name is added or moved to a host, the host object is added to the group.
7. Click **Preview** in the lower-left corner of the workspace, and verify that the hosts on which your virtual machines that include acct in the object name appear in the **Preview Group** window.
8. Click **Close**.
9. Click **Add another criteria set**.
   A new criteria set is added with the OR operator between the two criteria sets.
10. From the **Select the Object Type that matches the following criteria** drop-down menu, expand **vCenter Adapter**, select **Virtual Machine**, and configure the dynamic group criteria.
    a) From the criteria drop-down menu, select **Properties**.
    b) From the **Pick a property** drop-down menu, expand **Configuration** and double-click **Name**.
    c) From the operator drop-down menu, select **contains**.
    d) In the **Property value** text box, enter `acct`.

    You created a dynamic group where virtual machine objects with acct in the object name are included in the group that depends on the presence of those virtual machines. If a virtual machine with acct in the name is added to your environment, it is added to the group.
11. Click **Preview** in the lower-left corner of the workspace, and verify that the virtual machines with acct in the object name are added to the list that also includes the host systems.
12. Click **Close**.
13. Click **OK**.
    The Accounting VMs and Hosts group is added to the Groups list.

You created a dynamic object group that changes as virtual machines with acct in their names are added, removed, and moved in your environment.

Create a policy that determines how VMware Aria OperationsVMware Cloud Foundation Operations uses the alert definition to monitor your environment. See Create a Policy for the Accounting Alert.

### Create a Policy for the Accounting Alert

To configure how VMware Aria OperationsVMware Cloud Foundation Operations evaluates the accounting alert definition in your environment, you configure a policy that determines behavior so that you can apply the policy to an object group. The policy limits the application of the alert definition to only the members of the selected object group.

- Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.
- Verify that you created a group of objects that you use to manage you accounting objects. See Create a Custom Accounting Department Group.

When an alert definition is created, it is added to the default policy and activated, ensuring that any alert definitions that you create are active in your environment. This alert definition is intended to meet the needs of the accounting department, so you deactivate it in the default policy and create a new policy to govern how the alert definition is evaluated in your environment, including which accounting virtual machines and related hosts to monitor.

1. From the left menu, click **Operations › Configurations**, and then click the **Policy Definition** tile.
2. Click **Add**.
3. Type a name similar to `Accounting Objects Alerts Policy` and provide a useful description similar to the following example.

```
This policy is configured to generate alerts when

Accounting VMs and Hosts group objects are above trended

 CPU or memory usage.
```

4. Select **Default Policy** from the **Start with** drop-down menu.
5. On the left, click **Customize Alert / Symptom Definitions** and deactivate all the alert definitions except the new Acct VM CPU early warning alert.
   a) In the Alert Definitions area, click **Actions** and select **Select All**.
   The alerts on the current page are selected.
   b) Click **Actions** and select **Deactivate**.
   The alerts indicate Deactivated in the State column.
   c) Repeat the process on each page of the alerts list.
   d) Select **Acct VM CPU early warning** in the list, click **Actions** and select **Activate**.
   The Acct VM CPU early warning alert is now activated.
6. On the left, click **Apply Policy to Groups** and select **Accounting VMs and Hosts**.
7. Click **Save**.

You created a policy where the accounting alert definition exists in a custom policy that is applied only to the virtual machines and hosts for the accounting department.

Create an email notification so that you learn about alerts even you when you are not actively monitoring VMware Aria OperationsVMware Cloud Foundation Operations. See Configure Notifications for the Department Alert.

### Configure Notifications for the Department Alert

To receive an email notification when the accounting alert is generated, rather than relying on your ability to generally monitor the accounting department objects in VMware Aria OperationsVMware Cloud Foundation Operations, you create notification rules.

- Verify that you completed the alert definition for this scenario. See Add Recommendations to the Alert Definition.
- Verify that standard email outbound alerts are configured in your system. See Add a Standard Email Plug-In for Outbound Alerts.

Creating an email notification when accounting alerts are triggered is an optional process, but it provides you with the alert even when you are not currently working in VMware Aria OperationsVMware Cloud Foundation Operations.

1. From the left menu, click **Operations › Configurations**, and then click the **Notifications** tile.
2. Click **Add** to add a notification rule.
3. Configure the communication options.
   a) In the **Name** text box, type a name similar to `Acct Dept VMs or Hosts Alerts`.
   b) From the **Select Plug-In Type** drop-down menu, select **StandardEmailPlugin**.
   c) From the **Select Instance** drop-down menu, select the standard email instance that is configured to send messages.
   d) In the **Recipient(s)** text box, type your email address and the addresses of other recipients responsible for the accounting department alerts. Use a semicolon between recipients.
   e) Leave the **Notify again** text box blank.

      If you do not provide a value, the email notice is sent only once. This alert is a Risk alert and is intended as an early warning rather than requiring an immediate response.

You configured the name of the notification when it is sent to you and the method that is used to send the message.

4. In the Filtering Criteria area, configure the accounting alert notification trigger.
   a) From the **Notification Trigger** drop-down menu, select **Alert Definition**.
   b) Click **Select Alert Definitions**.
   c) Select **Acct VM CPU early warning** and click **Select**.
5. Click **Save**.

You created a notification rule that sends you and your designated engineers an email message when this alert is generated for your accounting department alert definition.

Create a dashboard with alert-related widgets so that you can monitor alerts for the accounting object group. See Create a Dashboard to Monitor Department Objects.

### Create a Dashboard to Monitor Department Objects

To monitor all the alerts related to the accounting department object group, you create a dashboard that includes the alert list and other widgets. The dashboard provides the alert data in a single location for all related objects.

Create an object group for the accounting department virtual machines and related objects. See Create a Custom Accounting Department Group.

Creating a dashboard to monitor the accounting virtual machines and related hosts is an optional process, but it provides you with a focused view of the accounting object group alerts and objects.

1. From the left menu, click **Operations › Dashboards**, and then click **Create**.
2. In the Dashboard Configuration definition area, type a tab name similar to `Accounting VMs and Hosts` and configure the layout options.
3. Click **Widget List** and drag the following widgets to the workspace.
   - **Alert List**
   - **Efficiency**
   - **Health**
   - **Risk**
   - **Top Alerts**
   - **Alert Volume**

   The blank widgets are added to the workspace. To change the order in which they appear, you can drag them to a different location in the workspace.
4. On the Alert List widget title bar, click **Edit Widget** and configure the settings.
   a) In the **Title** text box, change the title to `Acct Dept Alert List`.
   b) For the **Refresh Content** option, select **On**.
   c) Type `Accounting` in the **Search** text box and click **Search**.

      The Accounting value corresponds to the name of the object group for the accounting department virtual machines and related hosts.

   d) In the filtered resource list, select the **Accounting VMs and Hosts** group.
      The Accounting VMs and Hosts group is identified in the Selected Resource text box.
   e) Click **OK**.

   The Acct Dept Alert List is now configured to display alerts for the Accounting VMs and Hosts group objects.
5. Click **Widget Interactions** and configure the following interactions.
   a) For Acct Dept Alert List, leave the selected resources blank.
   b) For Top Alerts, Health, Risk, Efficiency, and Alert Volume select **Acct Dept Alert List** from the **Selected Resources** drop-down menu.

c) Click **Apply Interactions**.

With the widget interaction configured in this way, the select alert in the Acct Dept Alert List is the source for the data in the other widgets. When you select an alert in the alert list, the Health, Risk, and Efficiency widgets display alerts for that object, Top Alerts displays the topic issues affecting the health of the object, and Alert Volume displays an alert trend chart.

6. Click **Save**.

You created a dashboard that displays the alerts related to the accounting virtual machines and hosts group, including the Risk alert you created.

## Actions in VMware Aria OperationsVMware Cloud Foundation Operations

Actions
Actions
Actions are the ability to update objects or read data about objects in monitored systems, and are commonly provided in VMware Aria OperationsVMware Cloud Foundation Operations as part of a solution. The actions added by solutions are available from the object Actions menu, list and view menus, including some dashboard widgets, and can be added to alert definition recommendations.

The possible actions include read actions and update actions.

The read actions retrieve data from the target objects.

The update actions modifies the target objects. For example, you can configure an alert definition to notify you when a virtual machine is experiencing memory issues. Add an action in the recommendations that runs the Set Memory for Virtual Machine action. This action increases the memory and resolves the likely cause of the alert.

To see or use the actions for your vCenter objects, you must activate actions in the vCenter Adapter for each monitored vCenter instance. Actions can only be viewed and accessed if you have the required permissions.

### Actions, Modified Objects, and Object Levels in VMware Aria OperationsVMware Cloud Foundation Operations Actions

The list of actions includes the name of the action, the objects that each one modifies, and the object levels at which you can run the action. You use this information to ensure that you correctly apply the actions as alert recommendations and when the actions are available in the **Actions** menu.

#### Actions and Modified Objects

VMware Aria OperationsVMware Cloud Foundation Operations actions make changes to objects in your managed vCenter instances.

When you grant a user access to actions in VMware Aria OperationsVMware Cloud Foundation Operations, that user can take the granted action on any object that VMware Aria OperationsVMware Cloud Foundation Operations manages.

#### Action Object Levels

The actions are available when you work with different object levels, but they modify only the specified object. If you are working at the cluster level and select **Power On VM**, all the virtual machines in the cluster for which you have access permission are available for you to run the action. If you are working at the virtual machine level, only the selected virtual machine is available.

| Action | Modified Object | Object Levels |
|---|---|---|
| Rebalance Container | Virtual Machines | • Data Center |

*Table continued on next page*

*Continued from previous page*

| Action | Modified Object | Object Levels |
|---|---|---|
| | | • Custom Data Center |
| Delete Idle VM | Virtual Machines | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Set DRS Automation | Cluster | • Clusters |
| Move VM | Virtual Machine | • Virtual Machines |
| Power Off VM | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Shut Down Guest OS for VM | Virtual Machine<br><br>VMware Tools must be installed and running on the target virtual machines to run this action. | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Power On VM | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Delete Powered Off VM | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Set Memory for VM<br><br>and<br><br>Set Memory for VM Power Off Allowed | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Set Memory Resources for VM | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Set CPU Count for VM<br><br>and<br><br>Set CPU Count for VM Power Off Allowed | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Set CPU Resources for VM | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Set CPU Count and Memory for VM<br><br>and<br><br>Set CPU Count and Memory for VM Power Off Allowed | Virtual Machine | • Clusters<br>• Host Systems<br>• Virtual Machines |

*Table continued on next page*

*Continued from previous page*

| Action | Modified Object | Object Levels |
|---|---|---|
| Delete Unused Snapshots for VM | Snapshot | • Clusters<br>• Host Systems<br>• Virtual Machines |
| Delete Unused Snapshots for Datastore | Snapshot | • Clusters<br>• Datastores<br>• Host Systems |
| Execute Script | Virtual Machine | • Virtual Machine |
| Get Top Processes | Virtual Machine | • Virtual Machine |
| Apply Guest User Mapping | vCenter Server | • vCenter Server<br>**NOTE**<br>This action is deprecated and will be removed in the next release. |
| Clear Guest User Mapping | vCenter Server | • vCenter Server<br>**NOTE**<br>This action is deprecated and will be removed in the next release. |
| Export Guest User Mapping | vCenter Server | • vCenter Server<br>**NOTE**<br>This action is deprecated and will be removed in the next release. |
| Configure Included Services | Service Discovery Adapter Instance | • Service Discovery Adapter Instance<br>**NOTE**<br>This action is deprecated and will be removed in the next release. |

**Viewing Actions List in VMware Aria OperationsVMware Cloud Foundation Operations**

Viewing Actions List
Viewing Actions List
Actions are the method you use to configuration changes on managed objects that you initiate from VMware Aria OperationsVMware Cloud Foundation Operations. These actions are available to add to alert recommendations.

Actions are defined to run on the target object from different object levels, allowing you to add actions as recommendations for alert definitions that are configured for different base objects. The Actions page is a list of actions available in your environment.

## Where You Find the Actions List

To view the available actions, from the left menu, click **Operations › Configurations**, and then click the **Actions** tile.

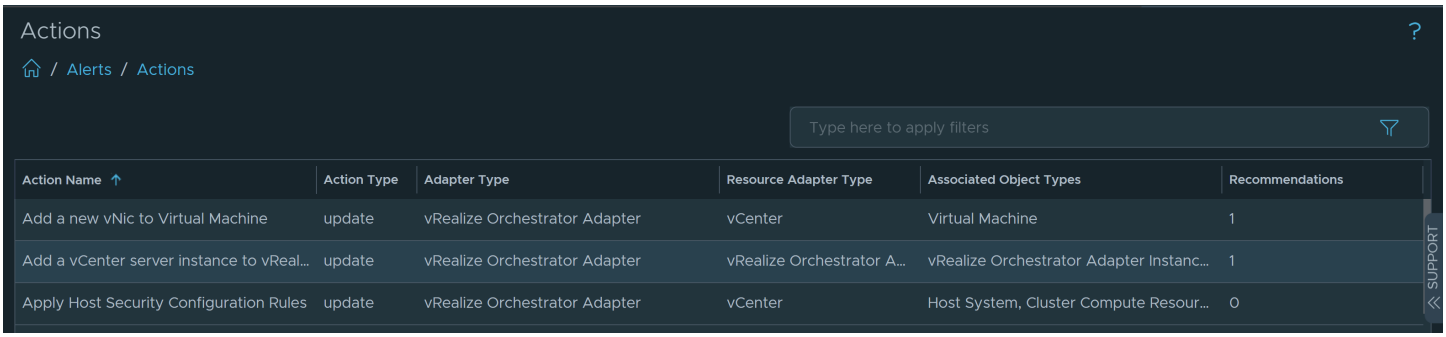


| Option | Description |
|---|---|
| Filter options | Limits the list based on the text you type.<br><br>You can also sort on the columns in the data grid. |
| Action Name | Name of the action. Duplicate names indicate that the action name is provided by more than one adapter or has more than one associated object.<br><br>Click this link to view Details page. On the Details Page, you can click on Recommendations to view the associated recommendations. |
| Action Type | Type of action that the action performs, either read or update.<br><br>• Update actions make changes to the target objects.<br>• Read actions retrieve data from the target objects. |
| Adapter Type | Name of the configured adapter that provides the action. |
| Resource Adapter Type | Adapter that provides the action. |
| Associated Object Types | Indicates the object level at which the action instance runs. |
| Recommendations | Indicates whether the action is used in at least one recommendation. |

The actions `Delete Unused Snapshots for Datastore Express` and `Delete Unused Snapshots for VM Express` appear. However, these actions can only be run in the user interface from an alert whose first recommendation is associated with this action. You can use the REST API to run these actions.

The following actions are also not visible except in the alert recommendations:

- `Set Memory for VM Power Off Allowed`
- `Set CPU Count for VM Power Off Allowed`
- `Set CPU Count and Memory for VM Power Off Allowed`

These actions are intended to be used to automate the actions with the `Power Off Allowed` flag set to true.

### Actions Supported for Automation

Recommendations can identify ways to remediate problems indicated by an alert. Some of these remediations can be associated with actions defined in your VMware Aria OperationsVMware Cloud Foundation Operations instance. You can automate several of these remediation actions for an alert when that recommendation is the first priority for that alert.

You activate actionable alerts in your policies. By default, automation is deactivated in policies. To configure automation for your policy, in the menu, click **Operations › Configurations**, and then click the **Policy Definition** tile. Then, to edit a

policy, access the **Alert / Symptom Definitions** workspace, and select **Local** for the **Automate** setting in the Alert / Symptom Definitions pane.

When an action is automated, you can use the **Automated** and **Alert** columns in **Recent Tasks** under **Administration › Control Panel** to identify the automated action and view the results of the action.

- VMware Aria OperationsVMware Cloud Foundation Operations uses the **automationAdmin** user account to trigger automated actions. For these automated actions that are triggered by alerts, the Submitted By column displays the **automationAdmin** user.
- The Alert column displays the alert that triggered the action. When an alert is triggered that is associated to the recommendation, it triggers the action without any user intervention.

The following actions are supported for automation:

- Delete Powered Off VM
- Delete Idle VM
- Move VM
- Power Off VM
- Power On VM
- Set CPU Count And Memory for VM
- Set CPU Count And Memory for VM Power Off Allowed
- Set CPU Count for VM
- Set CPU Count for VM Power Off Allowed
- Set CPU Resources for VM
- Set Memory for VM
- Set Memory for VM Power Off Allowed
- Set Memory Resources for VM
- Shut Down Guest OS for VM

## Roles Needed to Automate Actions

To automate actions, your role must have the following permissions:

- Create, edit, and import policies in **Operations › Configurations › Policy Definition**.
- Create, clone, edit, and import alert definitions in **Operations › Configurations › Alert Definitions**.
- Create, edit, and import recommendation definitions in **Operations › Configurations › Recommendations**.

> **IMPORTANT**
> You set the permissions used to run the actions separately from the alert and recommendation definition. Anyone who can modify alerts, recommendations, and policies can also automate the action, even if they do not have permission to run the action.

For example, if you do not have access to the Power Off VM action, but you can create and modify alerts and recommendations, you can see the Power Off VM action and assign it to an alert recommendation. Then, if you automate the action in your policy, VMware Aria OperationsVMware Cloud Foundation Operations uses the `automationAdmin` user to run the action.

## Example Action Supported for Automation

For the Alert Definition named `Virtual machine has chronic high CPU workload leading to CPU stress`, you can automate the action named `Set CPU Count for VM`.

When CPU stress on your virtual machines exceeds a critical, immediate, or warning level, the alert triggers the recommended action without user intervention.

### Integration of Actions with VMware Aria Automation

VMware Aria OperationsVMware Cloud Foundation Operations restricts actions on Datacentres and Custom Datacentres that contains VMware Aria Automation managed child objects such as, cluster compute resources, hosts, and VMs.

You can turn on or turn off the actions on VMware Aria Automation managed objects by modifying the **Operational Actions** from the respective vCentre in Cloud Accounts or by creating a new role with limited action ability on VMware Aria Automation managed objects.

#### Actions Determine Whether Objects Are Managed

Actions check the objects in the VMware Aria Automation managed resource container to determine which objects are being managed by VMware Aria Automation.

Actions such as Rebalance Container check the child objects of the data center container or custom data center container to determine whether the objects are managed by VMware Aria Automation. If the objects are being managed, the action does not appear on those objects.

### Working with Actions That Use Power Off Allowed

Some of the actions provided with VMware Aria OperationsVMware Cloud Foundation Operations require the virtual machines to shut down or power off, depending on the configuration of the target machines, to run the actions. You should understand the impact of the Power Off Allowed option before running the actions so that you select the best options for your target virtual machines.

#### Power Off and Shut Down

The actions that you can run on your vCenter instances include actions that shut down virtual machines and actions that power off virtual machines. It also includes actions where the virtual machine must be in a powered off state to complete the action. Whether the VM is shut down or powered off depends on how it is configured and what options you select when you run the action.

The shut-down action shuts down the guest operating system and then powers off the virtual machine. To shut down a virtual machine from VMware Aria OperationsVMware Cloud Foundation Operations, the VMware Tools must be installed and running on the target objects.

The power off action turns off the VM without regard for the state of the guest operating system. In this case, if the VM is running applications, your user might lose data. After the action is finished, for example, modifying the CPU count, the virtual machine is returned to the power state it was in when the action began.

#### Power Off Allowed and VMware Tools

For the actions where you are increasing the CPU count or the amount of memory on a VM, some operating systems support the actions if the Hot Plug is configured on the VM. For other operating systems, the virtual machine must be in a powered off state to change the configuration. To accommodate this need where the VMware Tools is not running, the Set CPU Count, Set Memory, and Set CPU Count and Memory actions include the Power Off Allowed option.

If you select Power Off Allowed, and the machine is running, the action verifies whether VMware Tools is installed and running.

- If VMware Tools is installed and running, the virtual machine is shut down before completing the action.
- If VMware Tools is not running or not installed, the virtual machine is powered off without regard for the state of the operating system.

If you do not select Power Off Allowed and you are decreasing the CPU count or memory, or the hot plug is not activated for increasing the CPU count or memory, the action does not run and the failure is reported in Recent Tasks.

**Power Off Allowed When Changing CPU Count or Memory**

When you run the actions that change the CPU count and the amount of memory, you must consider several factors to determine if you want to use the Power Off Allowed option. These factors include whether you are increasing or decreasing the CPU or memory and whether the target virtual machines are powered on. If you increase the CPU or memory values, whether hot plug is activated also affects how you apply the option when you run the action.

How you use Power Off Allowed when you are decreasing the CPU count or the amount of memory depends on the power state of the target virtual machines.

**Table 161: Decreasing CPU Count and Memory Behavior Based On Options**

| Virtual Machine Power State | Power Off Allowed Selected | Results |
|---|---|---|
| On | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, decreases the CPU or memory, and powers the machine back on.<br><br>If VMware Tools is not installed, the action powers off the virtual machine, decreases the CPU or memory, and powers the machine back on. |
| On | No | The action does not run on the virtual machine. |
| Off | Not applicable. The virtual machine is powered off. | The action decreases the value and leaves the virtual machine in a powered off state. |

How you use Power Off Allowed when you are increasing the CPU count or the amount of memory depends on several factors, including the state of the target virtual machine and whether hot plug is activated. Use the following information to determine which scenario applies to your target objects.

If you are increasing the CPU count, you must consider the power state of the virtual machine and whether CPU Hot Plug is activated when determining whether to apply Power Off Allowed.

**Table 162: Increasing CPU Count Behavior.**

| Virtual Machine Power State | CPU Hot Plug Activated | Power Off Allowed Selected | Results |
|---|---|---|---|
| On | Yes | No | The action increases the CPU count to the specified amount. |
| On | No | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, increases the CPU count, and powers the machine back on.<br><br>If VMware Tools is not installed, the action powers off the virtual machine, increases the CPU count, |

*Table continued on next page*

*Continued from previous page*

| Virtual Machine Power State | CPU Hot Plug Activated | Power Off Allowed Selected | Results |
|---|---|---|---|
| | | | and powers the machine back on. |
| Off | Not applicable. The virtual machine is powered off. | Not required. | The action increases the CPU count to the specified amount. |

If you are increasing the memory, you must consider the power state of the virtual machine, whether Memory Hot Plug is activated, and whether there is a Hot Memory Limit when determining how to apply Power Off Allowed.

**Table 163: Increasing Memory Amount Behavior**

| Virtual Machine Power State | Memory Hot Plug Activated | Hot Memory Limit | Power Off Allowed Selected | Results |
|---|---|---|---|---|
| On | Yes | New memory value ≤ hot memory limit | No | The action increases the memory the specified amount. |
| On | Yes | New memory value > hot memory limit | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on.<br><br>If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on. |
| On | No | Not applicable. The hot plug is not activated. | Yes | If VMware Tools is installed and running, the action shuts down the virtual machine, increases the memory, and powers the machine back on.<br><br>If VMware Tools is not installed, the action powers off the virtual machine, increases the memory, and powers the machine back on. |
| Off | Not applicable. The virtual machine is powered off. | Not applicable. | Not required | The action increases the memory the specified amount. |

## Configuring Objects

Using the power of object management - including metrics and alerts - you can monitor objects, applications, and systems that must stay up and running. Some metrics and alerts are prepackaged into dashboards and policies; others you combine into custom tools

VMware Aria OperationsVMware Cloud Foundation Operations discovers objects in your environment and makes them available to you. With the information that VMware Aria OperationsVMware Cloud Foundation Operations provides, you can quickly access and configure any object. For example, you can determine if a datastore is connected or providing data, or you can power on a virtual machine.

## Object Discovery

Its ability to monitor and collect data on objects in your systems environment makes VMware Aria OperationsVMware Cloud Foundation Operations a critical tool in maintaining system uptime and ensuring ongoing good health for all system resources from virtual machines to applications to storage - across physical, virtual, and cloud infrastructures.

Following are examples of objects that can be monitored.

- vCenter Server
- Virtual machines
- Servers/hosts
- Compute resources
- Resource pools
- Data centers
- Storage components
- Switches
- Port groups
- Datastores

**Adapters – Key to Object Discovery**

VMware Aria OperationsVMware Cloud Foundation Operations collects data and metrics from objects using adapters, that are the central components of management packs. You can customize adapter instances for your virtual environment using cloud accounts and other accounts. VMware Aria OperationsVMware Cloud Foundation Operations uses cloud accounts to manage the communication and integration with other products, applications, and functions.

- Cloud Accounts - You can configure cloud adapter instances and collect data from cloud solutions that are already installed in your cloud environment from the cloud accounts page.
- Other Accounts - You can view and configure native management packs and other solutions that are already installed and configure adapter instances from the other accounts page.
- Repository - You can activate or deactivate native management packs and add or upgrade other management packs from the Repository page.

The screenshot displays the list of available solutions in VMware Aria OperationsVMware Cloud Foundation Operations . You must first Activate the solution before adding and configuring the accounts.

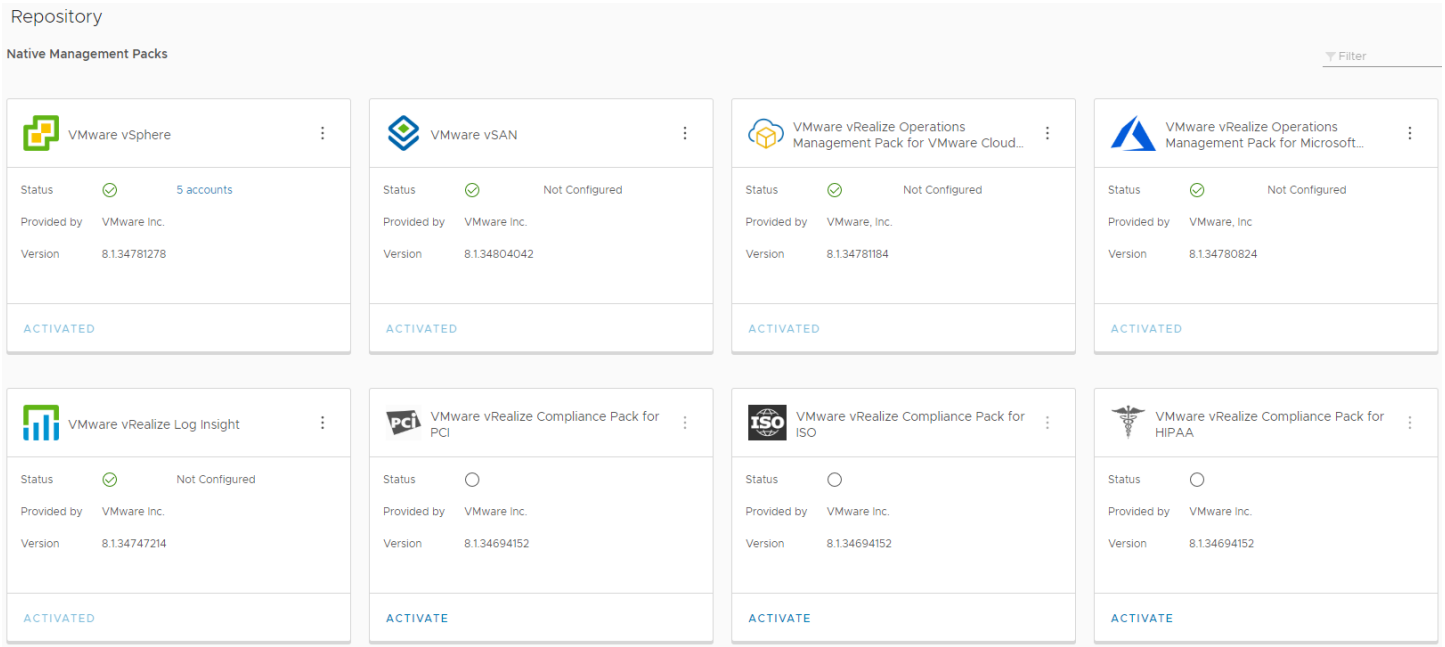


For complete information on configuring management packs and adapters, see Integrating Data Sources with VMware Aria OperationsVMware Cloud Foundation Operations

When you create a new adapter instance, it begins discovering and collecting data from the objects designated by the adapter, and notes the relationships between them. Now you can begin to manage your objects.

## Workload Management Inventory Objects

VMware Cloud Foundation OperationsVMware Cloud Foundation Operations discovers the following workload management objects and their child objects using the vCenter adapter:

- Tanzu Kubernetes cluster
- vSphere Pods
- Namespace

Clusters activated with vSphere IaaS control plane are called Supervisor Clusters. Kubernetes control plane is created inside the hypervisor layer. This layer contains specific objects that enable the capability to run Kubernetes workloads within ESXi. After a Supervisor is created, vSphere Namespaces can be created within the Supervisor. A vSphere Namespace sets the resource boundaries where vSphere Pods, VMs, and Tanzu Kubernetes Grid clusters can run.

> **NOTE**
> If a Supervisor Cluster spans multiple Compute Clusters, a Namespaces resource will be created for each Compute Cluster Resource and related to the Supervisor Cluster.

To understand the vSphere Tanzu Kubernetes architecture, see the *VMware vSphere with Tanzu* product documentation: https://docs.vmware.com/en/VMware-vSphere-with-Tanzu/index.html.

Workload management objects are excluded from the following workflows:

- Compliance
- Reclaim
- Rightsizing
- Workload optimization

## About Objects

Objects are the structural components of your mission-critical IT applications: virtual machines, datastores, virtual switches and port groups are examples of objects.

Because downtime equals cost - in unused resources and lost business opportunities - it's crucial that you successfully identify, monitor and track objects in your environment. The goal is to proactively isolate, troubleshoot and correct problems even before users are aware that anything is wrong.

When a user actually reports an issue, the solution should be quick and comprehensive.

For a complete list of objects that can be defined in VMware Aria OperationsVMware Cloud Foundation Operations refer to Object Discovery.

VMware Aria OperationsVMware Cloud Foundation Operations gives you visibility into objects including applications, storage and networks across physical, virtual and cloud infrastructures through a single interface that relates performance information to positive or negative events in the environment.

### Managing Objects

When you monitor a large infrastructure, the number of objects and corresponding metrics in VMware Aria OperationsVMware Cloud Foundation Operations grows rapidly, especially as you add solutions that extend dynamic monitoring and alerts to more parts of your infrastructure. VMware Aria OperationsVMware Cloud Foundation Operations gives you ample tools to stay abreast of events and issues.

### Adding Objects and Configuring Object Relationships

VMware Aria OperationsVMware Cloud Foundation Operations automatically discovers objects and their relationships once you create an adapter instance. You have the added ability to manually add any objects that you want monitored and to configure object relationships using abstract concepts rather than the connections recorded by VMware Aria OperationsVMware Cloud Foundation Operations. Where VMware Aria OperationsVMware Cloud Foundation Operations might discover the classic parent-child relationships between objects, you can create relationships between objects that might not normally be related. For example, you could configure all the datastores supporting a company department to be related.

When objects are related, a problem with one object appears as an anomaly on related objects. So object relationships can help you to identify problems in your environment quickly. The object relationships that you create are called custom groups.

### Custom Groups

To create an automated management system you need some way to organize objects so that you can quickly gain insights. You can achieve a high level of automation using custom groups. You have multiple options for tailoring group attributes to support your monitoring strategy.

For example, you can designate a group either to be static or to be updated automatically with membership criteria that you designate. Consider a non-static group of all virtual machines that are powered on and have OS type Linux. When you power on a new Linux VM, it is automatically added to the group and the policy is applied.

For additional flexibility, you can also specify individual objects to be always included or excluded from a given custom group. Or you can have a different set of alerts and capacity calculations for your production environment versus your testing environments.

### Managing Applications

VMware Aria OperationsVMware Cloud Foundation Operations allows you to create containers or objects that can contain a group of virtual machines or other objects in different structural tiers. This new application can then be managed as a single object, and have health badges and alarms aggregated from the child objects of the group.

For example, the system administrator of an online training system might request that you monitor components in the Web, application and database tiers of the training environment. You build an application that groups related training objects together in each tier. If a problem occurs with one of the objects, it is highlighted in the application display and you can investigate the source of the problem

## The Power of Object Management

Using the power of object management, including metrics and alerts - some prepackaged into dashboards and policies, others that you combine into custom monitoring tools - you'll keep a close watch on the objects, applications and systems that must stay up and running.

### Managing Objects in Your Environment

An object is the individual managed item in your environment for which VMware Aria OperationsVMware Cloud Foundation Operations collects data, such as a router, switch, database, virtual machine, host, and vCenter instances.

The system requires specific information about each object. When you configure an adapter instance, VMware Aria OperationsVMware Cloud Foundation Operations performs object discovery to start collecting data from the objects with which the adapter communicates.

An object can be a single entity, such as a database, or a container that holds other objects. For example, if you have multiple Web servers, you can define a single object for each Web server and define a separate container object to hold all of the Web server objects. Groups and applications are types of containers.

Categorize your objects using tags, so that you can easily find, group, or filter them later. A tag type can have multiple tag values. You or VMware Aria OperationsVMware Cloud Foundation Operations assigns objects to tag values. When you select a tag value, VMware Aria OperationsVMware Cloud Foundation Operations displays the objects associated with that tag. For example, if a tag type is Lifecycle and tag values are Development, Test, Pre-production, and Production, you might assign virtual machine objects VM1, VM2, or VM3 in your environment to one or more of these tag values, depending on the virtual machine function.

### Adding an Object to Your Environment

You might want to add an object by providing its information to VMware Aria OperationsVMware Cloud Foundation Operations. For example, some solutions cannot discover all the objects that might be monitored. For these solutions, you must either use manual discovery or manually add the object.

Verify that an adapter is present for the object you plan to add. See Connecting VMware Aria OperationsVMware Cloud Foundation Operations to Data Sources.

Verify that an adapter is present for the object you plan to add. See the Getting Started with *VMware Aria OperationsVMware Cloud Foundation Operations*.

> **NOTE**
> Objects added to VMware Aria OperationsVMware Cloud Foundation Operations via API will require an OSI license per object.

When you add an individual object, you provide specific information about it, including the kind of adapter to use to make the connection and the connection method. For example, a vSAN adapter does not know the location of the vSAN devices that you want to monitor.

1. From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile.
2. From the **Objects** tab, on the toolbar, click the plus sign.
3. Use the topic menus to reveal all fields and provide the required information.

| Option | Description |
|---|---|
| Display name | Enter a name for the object. For example, enter `vSAN-Host1`. |
| Description | Enter any description. For example, enter `vSAN-Host monitored with vSAN adapter` |
| Adapter type | Select an adapter type. For example, select **vSAN Adapter**. |
| Adapter instance | Select an adapter instance. |
| Object type | Select an object type. For a vSAN adapter, you might select vSAN-Host. When you select the object type, the dialog box selections change to include information you provide so that VMware Aria OperationsVMware Cloud Foundation Operations can find and connect with the selected object type. |
| Host IP address | Enter the host IP. For example, enter the IP address of vSAN-Host1. |
| Port number | Accept the default port number or enter a new value. |
| Credential | Select the Credential, or click the plus sign to add new login credentials for the object. |
| Collection interval | Enter the collection interval, in minutes. For example, if you expect the host to generate performance data every 5 minutes, set the collection interval to 5 minutes. |
| Dynamic Thresholding. | Accept the default, Yes. |

4. Click **OK** to add the object.

vSAN-Host1 appears in the Inventory as a host object type for the vSAN adapter type.

When you add an individual object, VMware Aria OperationsVMware Cloud Foundation Operations does not begin collecting metrics for the object until you turn on data collection. See Inventory : List of Objects.

For each new object, VMware Aria OperationsVMware Cloud Foundation Operations assigns tag values for its collector and its object type. Sometimes, you might want to assign other tags. See Creating and Assigning Tags.

For each new object, VMware Aria OperationsVMware Cloud Foundation Operations assigns tag values for its collector and its object type. Sometimes, you might want to assign other tags.

### Configuring Object Relationships

VMware Aria OperationsVMware Cloud Foundation Operations shows the relationship between objects in your environment. Most relationships are automatically formed when the objects are discovered by an installed adapter. In addition, you can use VMware Aria OperationsVMware Cloud Foundation Operations to create relationships between objects that might not normally be related.

Objects are related physically, logically, or structurally.

- Physical relationships represent how objects connect in the physical world. For example, virtual machines running on a host are physically related.
- Logical relationships represent business silos. For example, all the storage objects in an environment are related to one another.

- Structural relationships represent a business value. For example, all the virtual machines that support a database are structurally related.

Solutions use adapters to monitor the objects in your environment so that physical relationship changes are reflected in VMware Aria OperationsVMware Cloud Foundation Operations. To maintain logical or structural relationships, you can use VMware Aria OperationsVMware Cloud Foundation Operations to define the object relationships. When objects are related, a problem with one object appears as an influence on related objects. So object relationships can help you to identify problems in your environment quickly.

Apart from the parent-child relationship, you can also define new relationships in VMware Aria OperationsVMware Cloud Foundation Operations. The relationship between objects in your environment can be one-to-many, many-to-one, or one-one, the relationship can be defined in horizontal , vertical, or diagonal levels.

### Creating and Assigning Tags

A large enterprise can have thousands of objects defined in VMware Aria OperationsVMware Cloud Foundation Operations. Creating object tags and tag values makes it easier to find objects and metrics. With object tags, you select the tag value assigned to an object and view the list of objects that are associated with that tag value.

A tag is a type of information, for example, Adapter Types. Adapter Types is a predefined tag. Tag values are individual instances of that type of information. For example, when the system discovers objects using the vCenter Adapter, it assigns all the objects to the vCenter Adapter tag value under the Adapter Types tag.

You can assign any number of objects to each tag value, and you can assign a single object to tag values under any number of tags. You typically look for an object by looking under its adapter type, its object type, and possibly other tags.

If an object tag is locked, you cannot add objects to it. VMware Aria OperationsVMware Cloud Foundation Operations maintains locked object tags.

### Predefined Object Tags

VMware Aria OperationsVMware Cloud Foundation Operations includes several predefined object tags. It creates values for most of these tags and assigns objects to the values.

For example, when you add an object, the system assigns it to the tag value for the collector it uses and the kind of object that it is. VMware Aria OperationsVMware Cloud Foundation Operations creates tag values if they do not already exist.

If a predefined tag has no values, there is no object of that tag type. For example, if no applications are defined, the applications tag has no tag values.

Each tag value appears with the number of objects that have that tag. Tag values that have no objects appear with the value zero. You cannot delete the predefined tags or tag values.

**Table 164: Predefined Tags**

| Tag | Description |
|---|---|
| Collectors (Full Set) | Each defined collector is a tag value. Each object is assigned to the tag value for the collector that it uses when you add the object to VMware Aria OperationsVMware Cloud Foundation Operations. The default collector is VMware Aria OperationsVMware Cloud Foundation Operations Collector. |

*Table continued on next page*

*Continued from previous page*

| Tag | Description |
|---|---|
| Applications (Full Set) | Each defined application is a tag value. When you add a tier to an application, or an object to a tier in an application, the tier is assigned to that tag value. |
| Maintenance Schedules (Full Set) | Each defined maintenance schedule is a tag value, and objects are assigned to the value when you give them a schedule by adding or editing them. |
| Adapter Types | Each adapter type is a tag value, and each object that uses that adapter type is given the tag value. |
| Adapter Instances | Each adapter instance is a tag value, and each object is assigned the tag value for the adapter instance or instances through which its metrics are collected. |
| Object Types | Each type of object is a tag value, and each object is assigned to the tag value for its type when you add the object. |
| Recently Added Objects | The last day, seven days, 10 days, and 30 days have tag values. Objects have this tag value as long as the tag value applies to them. |
| Object Statuses | Tag value assigned to objects that are not receiving data. |
| Collection States | Tag value assigned to indicate the object collection state, such as collecting or not collecting. |
| Health Ranges | Good (green), Warning (yellow), Immediate (orange), Critical (red), and Unknown (blue) health statuses have tag values. Each object is assigned the value for its current health status. |
| Entire Enterprise | The only tag value is Entire Enterprise Applications. This tag value is assigned to each application. |
| Licensing | Tag values are License Groups found under **Subscriptions › Legacy Licenses › License Groups**. Objects are assigned to the license groups during VMware Aria OperationsVMware Cloud Foundation Operations installation. |
| Untag | Drag an object to this tag to delete the tag assignment. |

**Add an Object Tag and Assign Objects to the Tag**

An object tag is a type of information, and a tag value is an individual instance of that type of information. If the predefined object tags do not meet your needs, you can create your own object tags to categorize and manage objects in your environment. For example, you can add a tag for cloud objects and add tag values for different cloud names. Then you can assign objects to the cloud name.

Become familiar with the predefined object tags.

1. From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile.
2. Click the **Manage Tags** icon above the list of tags.
3. Click the **Add New Tag** icon to add a new row and type the name of the tag in the row.
   For example, type `Cloud Objects` and click **Update**.
4. With the new tag selected, click the **Add New Tag Value** icon to add a new row and type the name of the value in the row.
   For example, type `Video Cloud` and click **Update**.

5. Click **OK** to add the tag.
6. Click the tag to which you want to add objects to display the list of object tag values.
   For example, click **Cloud Objects** to display the Video Cloud object tag value.
7. Drag objects from the list in the right pane of the Inventory onto the tag value name.

   You can press Ctrl+click to select multiple individual objects or Shift+click to select a range of objects.

   For example, if you want to assign data centers that are connected through the vCenter Adapter, type `vCenter` in the search filter and select the data center objects to add.

### Use a Tag to Find an Object

The quickest way to find an object in VMware Aria OperationsVMware Cloud Foundation Operations is to use tags. Using tags is more efficient than searching through the entire object list.

Tag values that can also be tags are Applications and Object Types. For example, the Object Types tag has values for each object that is in VMware Aria OperationsVMware Cloud Foundation Operations, such as Virtual Machine, which includes all the virtual machine objects in your environment. Each of these virtual machines is also a tag value for the Virtual Machine tag. You can expand the tag value list to select the value for which you want to see objects.

1. From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile.
2. In the tag list in the center pane, click a tag for an object with an assigned value.

   When you click a tag, the list of values expands under the tag. The number of objects that is associated with each value appears next to the tag value.
   A plus sign next to a tag value indicates that the value is also a tag and that it contains other tag values. You can click the plus sign to see the subvalues.

3. Select the tag value.

   The objects that have that tag value appear in the pane on the right. If you select multiple tag values, the objects in the list depend on the values that you select.

| Tag Value Selection | Objects Displayed |
|---|---|
| More than one value for the same tag | The list includes objects that have either value. For example, if you select two values of the Object Types tag, such as Data Center and Host System, the list shows objects that have either value. |
| Values for two or more different tags | The list includes only objects that have all of the selected values. For example, if you select two values of the Object Types tag, such as Data Center and Host System, and you also select an adapter instance such as vC-1 of the vCenter Adapter instance tag, only Data Center or Host System objects associated with vC-1 appear in the list. Data Center or Host System objects associated with other adapter instances do not appear in the list, nor do objects that are not Data Center or Host System objects. |

4. Select the object from the list.

### Manage Object Tags Workspace

Manage Object Tags
Manage Object Tags
A large enterprise can have thousands of objects. When objects are assigned to a tag, and you choose to display objects with that tag value, the objects are easier to find on the Inventory list.

### Where You Find Manage Object Tags

From the left menu, click **Operations › Configurations** and then click the **Inventory Management** tile.

Click the **Manage Tags** icon above the list of tags in the middle pane.

### Manage Object Tags Options

The Manage Object Tags screen appears with previously created tags listed. In the left pane, you add tags. In the right pane, you add tag values.

- Click **Add a New Tag** and type a new tag name, or select a tag to delete.
- For the selected tag, click **Add a New Tag Value** and type a new tag value name, or select a tag value to delete.
- For the GEO Location tag, tag values are identified with a location on a world map. Select the tag value and click **Manage Location** to display the **Manage Location** map and pick a geographical location. Objects assigned to that tag value appear in that geographical location on the Inventory : Geographical Map of Objects.

## Manage Object Type Tags Workspace

Manage Object Type Tags
Manage Object Type Tags
Every object in your environment is of a particular object type. You use Manage Object Type Tags to control the object type tags displayed.

### How Manage Object Type Tags Works

For every adapter instance installed, VMware Aria OperationsVMware Cloud Foundation Operations discovers objects in your environment and starts collecting data from those objects.

### Where You Find Manage Object Type Tags

From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile. Click the **Manage Object Type Tags** icon above the list of tags.

### Manage Object Type Tags Options

Depending on the number of adapters installed , there may be hundreds of object type tags. The Manage Object Type Tags options allow you to turn on or off the tags listed.

- Type a filter word to show the object type tags with the word.
- Name lists all the object type tags.
- To toggle the display of an object type tag, select the check box in the Show Tag column of its row.

## Inventory : List of Objects

VMware Aria OperationsVMware Cloud Foundation Operations discovers objects in your environment for each adapter instance and lists them. From the complete list of all the objects in your environment, you can quickly access and configure any object. For example, you can check if a datastore is connected or providing data, or you can power on a virtual machine.

### How the List Works

Objects appear in a data grid. To find a particular object, you can sort a column in the grid or search for a filter word. In addition to sorting and searching, assigning objects to object tags makes it easier to find objects and metrics.

### Where You Find the List

From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile. The system lists all the objects in your environment.

### Inventory List Options

The center pane includes object tag options. The **Object** tab in the right pane includes toolbar options for all of the objects in your environment.

**Table 165: Object Tag Toolbar Options from the Center Pane**

| Option | Description |
|---|---|
| Collapse all | Closes all the tag group selections. |
| Deselect All | Tags remain selected until deselected. Use this option to deselect all tags. |
| Manage Tags | Add a tag or tag value. See Manage Object Tags Workspace. |
| Manage Object Type Tags | There might be many object type tags. Use this option to choose the object type tags to display. See Manage Object Type Tags Workspace. |

**NOTE**
If there are more than ten objects under an object tag in the centre pane, you can search for a descendant object using the Search option. If there are more than five hundred objects in a section, use the **View More** button under the last object that is displayed, to view the rest of the objects. Use the **Type to see more** button to find objects if there are more than one thousand objects.

Use the toolbar options to manage objects from the **Objects** Tab.

- Filter options limit the list to objects matching the filter. Filter options include ID, Name, Description, Maintenance Schedule, Adapter Type, Object Type, and Identifiers.
- Select the object to manage from the list. If an object tag is selected, only objects of the selected tag value are listed. Column headings help you to identify the object. See Object List Widget.

**Table 166: Inventory Toolbar Options from the Object Tab**

| Option | Description |
|---|---|
| Action | Perform an action on the selected object. Available actions depend on the object type. For example, Power on VM applies to the selected virtual machine. See Actions in VMware Aria Operations |
| Open in external application | If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application. For example, Open |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Virtual Machine in a vSphere Client or Search for VM logs in VMware Cloud Foundation Operations for logs. |
| Edit object | Edit the selected object. For example, add or change the maintenance schedule for a virtual machine. If multiple objects of the same type are selected, common identifiers for the object type are editable. For example, change the VM entity name of multiple datastores with a single edit. See Manage Objects Workspace. |
| Add object | VMware Aria OperationsVMware Cloud Foundation Operations discovers objects for most adapters. For adapters that do not support autodiscovery for all objects, the objects are manually added. See Manage Objects Workspace. |
| Discover Objects | Perform an IP scan to discover objects associated with a particular adapter. See Discover Objects Workspace. |
| Delete object | Remove the object from the list. |
| Start maintenance | Take the object offline for maintenance. See Manage Maintenance Schedules for Your Object Workspace. |
| End maintenance | Terminate the maintenance period and put the selected object back online. |
| Clear Selections | Clear all object selections. |
| Select All | Select all objects displayed. |
| Show Detail | Display the **Summary** tab of the selected object. |
| Per page | The number of objects to list per page. |

### Manage Objects Workspace

Manage Objects
Manage Objects
To collect data from an object, you might need to add an object or edit an existing object in your environment. For example, you might need to add objects for an adapter that does not support autodiscovery, or change the maintenance schedule of an existing object.

Where You Find Manage Objects
From the left menu, click **Environment** and then, click **Inventory**. Click the plus sign to add an object or the edit icon to edit the selected object.

Items that appear in the window depend on the object that you are editing. Not all options can be changed.

**Table 167: Manage Objects Add or Edit Options**

| Options | Description |
|---|---|
| Display name | Name of the object. Use only letters and numbers. Do not use nonalphanumeric characters or spaces. |
| Description | (Optional) For informational purposes only. |
| Adapter Type | If you are editing an object, you cannot change the adapter type. |
| Adapter Instance | If you are editing an object, you cannot change the adapter instance. |
| Object Type | If you are editing an object, you cannot change the object type. More configuration options might appear, depending on the object type. |
| Collection Interval | The collection interval for an object influences the collection status for the object. The collection interval for the adapter instance determines how often to collect data.<br><br>For example, if the collection interval for an adapter instance is set to five minutes, setting the collection interval for an object to 30 minutes prevents the object from having the No Data Receiving collection status after five collection cycles or 25 minutes.<br><br>The default value is 5 for adapter instances. You can increase this value, but not decrease it. To decrease it, contact VMware.<br><br>In cases of adapter instances such as VMwareAriaOperationsMgrAPI and HttpPost that push data to VMware Aria OperationsVMware Cloud Foundation Operations through the REST API, when data is no longer pushed, the status of the adapter instance is changed to Down after five collection intervals. For example, if the process pushes data every ten minutes and is stopped, the status of the adapter instance is changed to Down after 50 minutes. This behavior is expected for these adapter instance types. |
| Dynamic Thresholding | On by default, to activate dynamic thresholding and early warning smart alerts. See Dynamic Thresholds |

**Discover Objects Workspace**

Discover Objects
Discover Objects
If VMware Aria OperationsVMware Cloud Foundation Operations does not discover objects after an adapter instance is configured, use manual discovery. Discovering objects is more efficient than adding objects individually.

> **NOTE**
> You use discovery to define objects for embedded adapters. VMware Aria OperationsVMware Cloud Foundation Operations discovers objects that use external adapters.

Where You Find Discover Objects

From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile. Click **Discover Objects** in the tool bar of the **Objects** tab.

Discover Objects

The Discoveries section of the `describe.xml` file for the adapter might include parameters for discovery information. The `describe.xml` file is in the `conf` sub folder of the adapter, for example `xyz_adapter3/conf/describe.xml`.

| Options | Description |
|---|---|
| Collector | Collector that VMware Aria OperationsVMware Cloud Foundation Operations uses to discover objects. Only the VMware Aria OperationsVMware Cloud Foundation Operations Collector is added during installation. |
| Adapter Type | Adapter type for the objects to discover. |
| Adapter Instance | Adapter instance of the selected adapter type. |
| Discovery Info | Selection depends on the adapter type. For example, for a vCenter adapter, the Discovery Info selection adds an option to discover objects of a particular object type. |
| Only New Objects | On by default, to omit objects that are already discovered. |

Object Type

Based on your selection of Adapter Type, Adapter Instance, and Discovery Info, the Object Type values change. The available adapter types are NSX-T, vCenter Adapter, vSAN Adapter.

NSX-T Object Types:
- Edge Cluster
- Firewall Section
- Groups
- Load Balancer Pool
- Load Balancer Service
- Load Balancer Virtual Server
- Logical Router
- Logical Switch
- Router Service
- Transport Node
- Transport Zone

vCenter Object Types:
- Cluster Compute Resource
- Datastore
- Datacenter
- Folder
- Host System
- Resource Pool
- Virtual Machine

vSAN Object Types:
- Cache Disk
- Capacity Disk
- vSAN Cluster
- vSAN Disk Group
- vSAN Fault Domain
- vSAN Witness Host

### Discovery Results List

When you use the Discover Objects feature to manually discover objects in your environment, VMware Aria OperationsVMware Cloud Foundation Operations lists the objects of the specified object type. You can choose the objects to monitor.

Where You Find Discovery Results
From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile. Click **Discover Objects** from the tool bar in the **Objects** tab.

After you make selections in the Discover Objects Workspace, click **OK**. With the default setting, VMware Aria OperationsVMware Cloud Foundation Operations displays only newly discovered objects. See Discover Objects Workspace.

**Table 168: Object Types**

| Options | Description |
|---|---|
| Object Type | Discovered object types of the Object Type selected on the Discover Objects Workspace. |
| Object Count | Number of objects of the object type. |
| Import | When selected, imports the object type. Option is active and selectable for newly discovered object types. |
| Collect | When selected, imports the object type and starts collecting data. Option is active and selectable for newly discovered object types. |
| Credential | If the object type requires a login credential to collect data from the object., the value is **True**. |

Double-click the Object Type to display a list of objects to monitor.

**Table 169: Objects**

| Options | Description |
|---|---|
| Object | Objects of the selected type that exist in the environment for the adapter. For example, the vCenter adapter discovers objects in the vCenter system. |
| Import | When selected, imports the object but does not start collecting data. Option is active and selectable for newly discovered objects that do not exist in the VMware Aria OperationsVMware Cloud Foundation Operations environment . |
| Exists | Indicates that the object exists in the VMware Aria OperationsVMware Cloud Foundation Operations environment. |
| Collect | When selected, imports the object and starts collecting data. Option is active and selectable for newly discovered objects that do not exist in the VMware Aria OperationsVMware Cloud Foundation Operations environment. |

## Manage Maintenance Schedules for Your Object Workspace

Manage Maintenance Schedules
Manage Maintenance Schedules
You use maintenance mode to take an object offline. Many objects in your environment might be intentionally taken offline. For example, you might deactivate a server to update software. If VMware Aria OperationsVMware Cloud Foundation Operations collects metrics when the object is offline, it might generate incorrect alerts that affect the data for the object's health. When an object is in maintenance mode, VMware Aria OperationsVMware Cloud Foundation Operations does not collect metrics from the object and does not generate alerts for it.

How Maintenance Schedules Work
If an object undergoes maintenance at fixed intervals, you can create a maintenance schedule and assign it to the object. For example, you can put an object into maintenance mode from midnight until 3 a.m. every Tuesday night. You can also manually put an object in maintenance mode, either indefinitely or for a specified period of time. These methods are not mutually exclusive. You can put an object in maintenance mode or take it out of maintenance mode, even if it has an assigned maintenance schedule.
Where You Find Manage Maintenance Schedules
From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile. Click **Start Maintenance** from the tool bar in the **Objects** tab.

**Table 170: Manage Maintenance Schedules Options**

| Options | Description |
|---|---|
| I will come back and end maintenance myself. | Maintenance mode starts for the selected object when you click **OK**. You must manually end maintenance mode for this object. |
| End maintenance in | Type the number of minutes that the object is in maintenance mode. |
| End maintenance on | Click the calendar icon, and select the date that maintenance mode ends. |

## Define Custom Property Workspace

In VMware Aria OperationsVMware Cloud Foundation Operations, you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign custom properties to any subset of objects irrespective of the adapter kind and resource kind. You can use a mouse click, search filter, or a tag selector to select the correct object.

Where You Find Add/Edit Custom Property
From the left menu, click **Operations › Configurations**, and then click **Inventory Management** tile. Click **Add/Edit Custom Property** from the tool bar in the **Objects** tab.

**Table 171: Add/Edit Custom Property**

| Options | Description |
|---|---|
| Property Name | Select or enter a property name. |
| Type | Select the property type from the drop-down menu. |
| Value | Enter a value for the property. |

You can assign the custom properties defined in this page to the Custom Object Groups and New Groups.

For more information, see Custom Object Groups Workspace to Create a New Group.

**Inventory : Geographical Map of Objects**

VMware Aria OperationsVMware Cloud Foundation Operations discovers objects in your environment for each adapter. Objects that are assigned a GEO Location tag appear on a geographical map. You can use this map to quickly locate your objects in the world.

**How the Geographical Map Works**

Objects with the GEO Location tag appear on a map of the world.

- To create a GEO Location tag, see Manage Object Tags Workspace.
- To assign objects to the tag, see Creating and Assigning Tags.

**Where You Find the Geographical Map**

From the left menu, click **Operations › Configurations**, and then click the **Inventory Management** tile. Click the **Geographical** tab.

**Geographical Map Options**

Use the plus sign to zoom in. Use the minus sign to zoom out. Click and drag to pan the map to the left or right.

**Managing Custom Object Groups in VMware Aria OperationsVMware Cloud Foundation Operations**

Managing Custom Object Groups
Managing Custom Object Groups
A custom object group is a container that includes one or more objects. VMware Aria OperationsVMware Cloud Foundation Operations uses custom groups to collect data from the objects in the group, and report on the data collected.

**Why Use Custom Object Groups?**

You use groups to categorize your objects and have the system collect data from the groups of objects and display the results in dashboards and views according to the way you define the data to appear.

You can create static groups of objects, or dynamic groups with criteria that determine group membership as VMware Aria OperationsVMware Cloud Foundation Operations discovers and collects data from new objects added to the environment.

VMware Aria OperationsVMware Cloud Foundation Operations provides commonly used object group types, such as World, Environment, and Licensing. The system uses the object group types to categorize groups of objects. You assign a group type to each group so that you can categorize and organize the groups of objects that you create.

**Types of Custom Object Groups**

When you create custom groups, you can use rules to apply dynamic membership of objects to the group, or you can manually add the objects to the group. When you add an adapter, the groups associated with the adapter become available in VMware Aria OperationsVMware Cloud Foundation Operations.

- Dynamic group membership. To dynamically update the membership of objects in a group, define rules when you create a group. VMware Aria OperationsVMware Cloud Foundation Operations adds objects to the group based on the criteria that you define.
- Mixed membership, which includes dynamic and manual.
- Manual group membership. From the inventory of objects, you select objects to add as members to the group.

- Groups associated with adapters. Each adapter manages the membership of the group. For example, the vCenter Server adapter adds groups such as datastore, host, and network, for the container objects in the vSphere inventory. To modify these groups, you must do so in the adapter.

Administrators of VMware Aria OperationsVMware Cloud Foundation Operations can set advanced permissions on custom groups. Users who have privileges to create groups can create custom groups of objects and have VMware Aria OperationsVMware Cloud Foundation Operations apply a policy to each group to collect data from the objects and report the results in dashboards and views.

When you create a custom group, and assign a policy to the group, the system uses the criteria defined in the applied policy to collect data from and analyze the objects in the group. VMware Aria OperationsVMware Cloud Foundation Operations reports on the status, problems, and recommendations for those objects based on the settings in the policy.

> **NOTE**
> Only custom groups defined explicitly by users can be exported from or imported to VMware Aria OperationsVMware Cloud Foundation Operations. Users are able to export or import multiple custom groups. Once an import function has been executed, the user must check to determine if a policy or policies should be associated with the imported group. Export-import operations are available for user defined (created explicitly by user) custom groups only.

## How Policies Help VMware Aria OperationsVMware Cloud Foundation Operations Report On Object Groups

When you apply a policy to an object group, VMware Aria OperationsVMware Cloud Foundation Operations uses threshold settings, metrics, super metrics, attributes, properties, alert definitions, and problem definitions that you activated in the policy to collect data from the objects in the group, and report the results in dashboards and views.

When you create a new object group, you have the option to apply a policy to the group.

- To associate a policy with the custom object group, select the policy in the group creation wizard.
- To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

VMware Aria OperationsVMware Cloud Foundation Operations applies policies in priority order, as they appear on the Active Policies tab. When you establish the priority for your policies, VMware Aria OperationsVMware Cloud Foundation Operations applies the configured settings in the policies according to the policy rank order to analyze and report on your objects. To change the priority of a policy, you click and drag a policy row. The default policy is always kept at the bottom of the priority list, and the remaining list of active policies starts at priority 1, which indicates the highest priority policy. When you assign an object to be a member of multiple object groups, and you assign a different policy to each object group, VMware Aria OperationsVMware Cloud Foundation Operations associates the highest ranking policy with that object.

### User Scenario: Creating Custom Object Groups

As a system administrator, you must monitor the capacity for your clusters, hosts, and virtual machines. VMware Aria OperationsVMware Cloud Foundation Operations monitors them at different service levels to ensure that these objects adhere to the policies established for your IT department, and discovers and monitors new objects added to the environment. You have VMware Aria OperationsVMware Cloud Foundation Operations apply policies to the object groups to analyze, monitor, and report on the status of their capacity levels.

- Know the objects that exist in your environment, and the service levels that they support.
- Understand the policies required to monitor your objects.
- Verify that policies are available to monitor the capacity of your objects.

To have VMware Aria OperationsVMware Cloud Foundation Operations monitor the capacity levels for your objects to ensure that they adhere to your policies for your service levels, you categorize your objects into Platinum, Gold, and Silver object groups to support the service tiers established.
You create a group type, and create dynamic object groups for each service level. You define membership criteria for each dynamic object group to have VMware Aria OperationsVMware Cloud Foundation Operations keep the membership of

objects current. For each dynamic object group, you assign the group type, and add criteria to maintain membership of your objects in the group. To associate a policy with the custom object group, you can select the policy in the group creation wizard.

1. To create a group type to identify service level monitoring, from the left menu click **Operations › Configurations**, and then click the **Custom Groups** tile. From the **Custom Groups** page, click **Group Types** from the top of the page.
2. On the Group Types toolbar, click **Add** and type `Service Level Capacity` for the group type.
   Your group type appears in the list.
3. From the left menu, click **Operations › Configurations**, and then click the **Custom Groups** tile.
4. To create a new object group, click **Add**.

   The New Group workspace appears where you define the data and membership criteria for the dynamic group.

   a) In the Name text box, enter a name for the object group, such as `Platinum_Objects`.
   b) In the **Group Type** drop-down menu, select **Service Level Capacity**.
   c) (Optional) In the **Policy** drop-down menu, select your service level policy that has thresholds set to monitor the capacity of your objects.
      To associate a policy with the custom object group, select the policy in the group creation wizard. To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.
   d) Select the **Keep group membership up to date** check box so that VMware Aria OperationsVMware Cloud Foundation Operations can discover objects that meet the criteria, and add those objects to the group.
5. Define the membership for virtual machines in your new dynamic object group to monitor them as platinum objects.
   a) From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Virtual Machine**.
   b) From the empty drop-down menu for the criteria, select **Metrics**.
   c) From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.
   d) From the conditional value drop-down menu, select **is less than**.
   e) From the **Metric value** drop-down menu, type `10`.
6. Define the membership for host systems in your new dynamic object group to monitor them as platinum objects.
   a) Click **Add another criteria set**.
   b) From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Host System**.
   c) From the empty drop-down menu for the criteria, select **Metrics**.
   d) From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **Current Size**.
   e) From the conditional value drop-down menu, select **is less than**.
   f) From the **Metric value** drop-down menu, type `100`.
7. Define the membership for cluster compute resources in your new dynamic object group.
   a) Click **Add another criteria set**.
   b) From the **Select Object** drop-down menu, select **vCenter Adapter**, and select **Cluster Compute Resources**.
   c) From the empty drop-down menu for the criteria, select **Metrics**.
   d) From the **Pick a metric** drop-down menu, select **Disk Space** and double-click **capacityRemaining**.
   e) From the conditional value drop-down menu, select **is less than**.
   f) From the **Metric value** drop-down menu, type `1000`.
   g) Click **Preview** to determine whether objects already match this criteria.
8. Click **OK** to save your group.
   When you save your new dynamic group, the group appears in the Service Level Capacity folder, and in the list of groups on the **Groups** tab.

9. Wait five minutes for VMware Aria OperationsVMware Cloud Foundation Operations to collect data from the objects in your environment.

VMware Aria OperationsVMware Cloud Foundation Operations collects data from the cluster compute resources, host systems, and virtual machines in your environment, according to the metrics that you defined in the group and the thresholds defined in the policy that is applied to the group, and displays the results about your objects in dashboards and views.

To monitor the capacity levels for your platinum objects, create a dashboard, and add widgets to the dashboard. See Dashboards in .

### Object Group Types in VMware Aria OperationsVMware Cloud Foundation Operations

Object Group Types
Object Group Types
An object group type is an identifier that you apply to a specific group of objects in your environment to categorize them. You can add new group types, and apply them to groups of objects so that VMware Aria OperationsVMware Cloud Foundation Operations can collect data from the object group and display the results in the dashboards and views.

#### How the Group Types Work

Use group types to categorize your objects so that the system can apply policies to them to track, and display specific status, such as alerts, workload, faults, risk, and so on.

When you create a new group type, VMware Aria OperationsVMware Cloud Foundation Operations adds it to the existing list of group types, and creates a new folder with the name of your group type in the Environment Custom Groups list.

When you create a new group of objects, you assign a group type to that group of objects. You add objects from the inventory trees to your custom group, then create your dashboard, add widgets to the dashboard, and configure the widgets to display the data collected from the objects in the group. You can then monitor and manage the objects.

You can apply a group type to a group of objects that you create manually, or to object groups that you cannot modify, such those added by adapters. Each adapter that you add to VMware Aria OperationsVMware Cloud Foundation Operations adds one or more static groups of objects to group the data received from the adapter sources.

The list of group types appears in the Content area under Group Types. The custom object groups appear in the Environment area under Custom Groups.

#### Where You Create and Modify a Group Type

From the left menu, click **Inventory** and then click **Custom Group and Datacenters**. Click **Group Types** next to custom groups. You can add, edit, delete, and select groups from the group types page.

#### Group Type Options

You can add, edit, or delete group types. You cannot edit group types that are created by adapters.

### Groups Tab on the Environment Overview Pane

Groups are containers that can contain any number and type of objects in your environment. VMware Aria OperationsVMware Cloud Foundation Operations collects data from the objects in the group and displays the results in dashboards and views that you define.

## How Groups Work

Groups are installed with VMware Aria OperationsVMware Cloud Foundation Operations, created by an adapter, or created by a user. Based on the group criteria, you can use groups to organize your environment and monitor all objects in the group together. You can also assign policies to groups and make group membership dynamic.

For example, if you have a set of vSphere hosts and you do not want to generate alerts when the host goes into maintenance mode, you can put the vSphere hosts in a group and assign a policy that includes a maintenance schedule setting. During the maintenance period, VMware Aria OperationsVMware Cloud Foundation Operations ignores any metrics for those objects and does not generate any alerts. After the maintenance period ends, VMware Aria OperationsVMware Cloud Foundation Operations returns to monitoring the objects and generates alerts if an outage occurs.

## Where You Find Custom Groups

From the left menu, click **Operations › Configurations**, and then click the **Custom Groups** tile.

## Custom Group Options

Click **ADD** button to add a group. You can only edit, clone, or delete a user-created group. You cannot modify groups installed with VMware Aria OperationsVMware Cloud Foundation Operations or by an adapter.

You can click the **Horizontal Ellipses** to import or export the custom group. The Groups data grid displays an overview of the state of each group. You can use the All Filters option to sort the custom groups based on Name, ID, Group Type, and Description columns.

To sort the list of custom groups based on columns, click the column heading of the following columns:

- Name
- Health
- Risk
- Efficiency
- Description
- Members Count

**Table 172: Group Data Grid Options**

| Option | Description |
|---|---|
| Name | Select the group name to display a summary of the group. Select to the right of the name to edit, clone, or delete the group. |
| Summary | Criticality of the health, risk, and efficiency of any group. Click a group with a red, orange, or yellow criticality to get more details about potential problems with objects in the group. |
| Members Count | Displays the number of members in the selected group. |
| Policy | Displays the policy associated with the selected group. |
| Dynamic Membership | Displays whether the group is static or dynamic. The available options are true and false. |
| Defined by | Displays who has defined the attributes of the group. The available options are:<br>• System<br>• User Defined<br>• Management Pack |

## Custom Object Groups Workspace

You can create and edit custom groups of objects to have VMware Aria OperationsVMware Cloud Foundation Operations collect data from the objects and display the results in the dashboards and views so that you can monitor your objects and take action on them when problems occur.

### How the Custom Groups Workspace Works

When you create a new object group, you define a meaningful group name, and select the group type. To associate the custom object group with a policy for analysis, you select the policy in the group creation wizard. You can leave the policy selection blank to not associate a policy with the object group. When the policy selection is blank, the custom object group is associated with the policy that is designated as the default policy.

You select the object types, and determine whether membership in the object group is static, dynamic, or a combination of static and dynamic membership.

- To create a static object group, you add objects to the group. You do not include criteria for object membership.
- To create a dynamic object group that VMware Aria OperationsVMware Cloud Foundation Operations updates based on specific criteria, you select the object type and define membership criteria for the group based on metrics, relationships, and properties.

When you add objects to a custom object group, a new folder appears in the Custom Groups navigation pane on the left, and includes the member objects.

### Where You Create and Modify Object Groups

To create or modify static or dynamic object groups, or object groups that have a combination of static and dynamic membership, click **Operations › Configurations**, and then click the **Custom Groups** tile. The **Custom Groups** tab displays a list of custom object groups, and the object groups for adapters added to VMware Aria OperationsVMware Cloud Foundation Operations.

To edit existing groups, select a group, and then click the vertical ellipsis and select **Edit**.

## Custom Object Groups Workspace to Create a New Group

You can create a new object group, define custom properties, assign a group type and objects to the group. When you create the group, you can assign a policy, or leave the policy selection blank to apply the default policy. VMware Aria OperationsVMware Cloud Foundation Operations collects data from the objects in the group based on the settings in the policy that is associated with the group. The results appear in the dashboards and views.

#### Where You Assign Custom Group Type, Policy, and Membership

To assign the group type, policy, and membership, click **Operations › Configurations**, and then click the **Custom Groups** tile. Click **Add** to add a new group. In the New Group workspace, you can define the membership criteria, and select the objects to include or exclude.

To associate a policy with the custom object group, select the policy in the group creation wizard. To not associate a specific policy with the object group, leave the policy selection blank. The custom object group will be associated with the default policy. If the default policy changes, this object group will be associated with the new default policy.

**Table 173: New Group Workspace**

| Option | Description |
|---|---|
| **Custom Group Tab** | |
| Name | Meaningful name of the object group. |
| | Meaningful description for the object group. |
| Group Type | Categorization for the object group. New custom groups appear in a dedicated folder in the Custom Groups navigation pane on the left. |
| Policy | Assigns a policy to one or more groups of objects to have VMware Aria OperationsVMware Cloud Foundation Operations analyze the objects according to the settings in your policy, trigger alerts when the defined thresholds are violated, and display the results in dashboards, views, and reports. You can assign a policy to the group when you create the group, or you can assign it later from the edit custom group wizard or from the policies area. |
| Keep group membership up to date | For dynamic object groups, VMware Aria OperationsVMware Cloud Foundation Operations can discover objects that match the criteria for the group membership according to the rules that you define, and update the group members based on the search results. |
| **Define Membership Criteria Tab** | |
| Membership Criteria | Defines the criteria for a dynamic object group and has VMware Aria OperationsVMware Cloud Foundation Operations keep the object membership of the group current.<br>• Object Type drop-down menu. Selects the type of objects to add to the group, such as virtual machines.<br>• Metrics, Relationship, and Properties criteria drop-down menu. Defines the criteria for VMware Aria OperationsVMware Cloud Foundation Operations to apply to collect data from the selected objects.<br>– Metrics. An instance of a data type, or attribute, that varies based on the object type. A metric is used as measurement criteria to collect data from objects. For example, you can select system attributes as a metric, where an attribute is a type of data that VMware Aria OperationsVMware Cloud Foundation Operations collects from objects.<br>– Relationship. Indicates how the object is related to other objects. For example, you can require a virtual machine object to be a child object that contains a certain word in the vSphere Hosts and Clusters navigation tree.<br>– Properties. Identifies a configuration parameter for the object. For example, you can require a virtual machine to have a memory limit that is greater than 100KB.<br>– Object Name. Specifies the name of the object. For example, you can create custom group which contains object with specific names.<br>– Tag. Specifies a tag for the object. For example, you can create membership criteria based on specific tags.<br>**NOTE**<br>For the Metrics and Properties criteria drop-down menu, the metrics presented by default are a subset of available metrics. If the desired metric is not represented, use the **Select Object** button in the title bar of the data selection tree to re-filter the displayed list.<br>• You can use the filtering option to define the membership criteria for a group. The available options are: |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | – is - Displays the result of the typed text. Example - "!", displays all results equal to the symbol "!"<br>– is not - Displays the result of the typed text. Example - "1_node", displays all results not equal to "1_node"<br>– contains -Displays the result of the typed text. Example - "An", displays all results that contain "An"<br>– does not contain - Displays the result of the typed text. Example - "An", displays all results that does not contain "An"<br>– starts with - Displays the result of the typed text. Example - "!!", displays all results that starts with "!!"<br>– ends with - Displays the result of the typed text. Example - "N", displays all results that ends with "N"<br>– does not start with - Displays the result of the typed text. Example - "N", displays all results that does not start with "N"<br>– does not end with - Displays the result of the typed text. Example - "S", displays all results that does not end with "S"<br>– matches regular expression - Displays the result if it matches the regular expression. Example .*8\.\d*<br>– does not match regular expression Displays the result if it does not match the regular expression. Example .*8\.\d*<br>• Add. Includes another metric, relationship, or property for the object type.<br>• Remove. Deletes the selected object type from the membership criteria, or delete the selected metric, relationship, or property type from the criteria for the object type.<br>• Add another criteria set. Adds another object type to add to the group. For example, you might want to create a single object group to track vCenter instances and Host Systems. |
| **Objects to Include/Exclude Tab** | |
| Objects To Always Exclude | Determine which objects to exclude from the group every time VMware Aria OperationsVMware Cloud Foundation Operations collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define for membership. In previous versions of VMware Aria OperationsVMwa re Cloud Foundation Operations, these objects were called a denylist.<br><br>• **Select Objects Pane > Inventory Tab**. Displays the list of available object groups. Click on an object group to view the objects in the right pane. You can select some or all the objects in the object group and drag and drop it to the left pane for exclusion from the custom group.<br>• **Select Objects Pane > Custom Group Tab**. Displays the list of available custom groups. Click on a custom group and drag and drop it to the left pane for exclusion from the custom group.<br>• From the left pane, select the objects you want to delete and click **Remove** to delete the objects that you have added for exclusion from the custom group. |
| Objects To Always Include | Determine which objects to include in the group every time VMware Aria OperationsV Mware Cloud Foundation Operations collects data from the objects, regardless of the membership criteria. The objects that you include override the criteria that you define |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | for membership. In previous versions of VMware Aria OperationsVMware Cloud Foundation Operations, these objects were called an allowlist.<br>• **Select Objects Pane > Inventory Tab**. Displays the list of available object groups. Click on an object group to view the objects in the right pane. You can select some or all the objects in the object group and drag and drop them to the left pane for inclusion into the custom group.<br>• **Select Objects Pane > Custom Group Tab**. Displays the list of available custom groups. Click on a custom group and drag and drop it to the left pane for inclusion into the custom group.<br>• From the left pane, select the objects you want to delete and click **Remove** to delete the objects that you have added for inclusion to the custom group. |
| **Assign Custom Properties Tab** | |
| Assign Custom Properties | In VMware Aria OperationsVMware Cloud Foundation Operations, you can define custom properties to collect and store operational data related to different objects. The custom property can be either a string or a numeric. You can assign the newly defined custom properties to new groups or existing groups.<br>• Property Name. Select or specify a name for the custom property.<br>• Type. Select the type of custom property from the drop-down menu. The custom property can either be a string or a numeric.<br>• Value. Specify a custom property value, which should be assigned to this custom property when an object is added to the group.<br>• Reset Value. Specify a custom property value, which should be assigned to this custom property when an object leaves the group.<br>• **Reset**. Resets the custom property to a non-zero value.<br>• **Remove**. Removes the custom property from the group.<br>• **Add Another Custom Property**. Adds another custom property to the group. |

- Use the **Previous** and **Next** buttons to navigate between the tabs.
- Use the **Next** button to move to the next tab.
- Use **Create** button to save and create the custom group.
- Use the **Preview** button to preview the list of objects in the group and to verify that the criteria you defined is applicable to the group of objects. If the criteria that you defined is valid, the preview displays applicable objects. If the criteria is not valid, the preview does not display any objects.

### Managing Application Groups

An application is a container construct that represents a collection of interdependent hardware and software components that deliver a specific capability to support your business. VMware Aria OperationsVMware Cloud Foundation Operations builds an application to determine how your environment is affected when one or more components in an application experiences problems, and to monitor the overall health and performance of the application. Object membership in an application is not dynamic. To change the application, you manually modify the objects in the container.

#### Reasons to Use Applications

VMware Aria OperationsVMware Cloud Foundation Operations collects data from components in the application and displays the results in a summary dashboard for each application with a real-time analysis for any of the components. If a component experiences problems, you can see where in the application the problems arise, and determine how problems spread to other objects.

**NOTE**
VMware Aria OperationsVMware Cloud Foundation Operations provides for calendar periodicity. If your application includes work performed on a specific day of the month, for example, the 15th of the month or the last day of the month, this calendar function identifies the pattern after six cycles of the application. Once the pattern is recognized, the system can forecast accurately into the future. Because the system acquires its information from the input data, you do not have to give any details about how you schedule periodical work.

### Applications Tab on the Environment Overview Pane

Applications are groups of related objects in your environment that mimic an application in your business. Use the summary to track the health of objects in the application and help troubleshoot performance issues.

#### How Applications Work

In VMware Aria OperationsVMware Cloud Foundation Operations, each application contains one or more tiers and each tier contains one or more objects. The tier is a convenient way to organize objects that perform a specific task in an application. For example, you can group all of your database servers together in a tier.

The objects in a tier are static. If the set of objects in a tier changes, you must manually edit the application.

Construct an application to view a particular segment of your business. The application shows how the performance of one object affects other objects in the same application, and helps you to locate the source of a problem. For example, if you have an application that includes all the database, Web, and network servers that process sales data for your business, you see a yellow, orange, or red status if the application health is degrading. Starting with the application summary dashboard, you can investigate which server is causing or exhibiting the problem.

#### Where You Find Applications

From the left menu, click **Operations › Applications**.

Applications defined in a previous release of VMware Aria OperationsVMware Cloud Foundation Operations appear after an upgrade.

#### Application Options

Select an application to edit or delete, or click the **ADD** button to add an application.

The Applications data grid displays an overview of the state of each application.

**Table 174: Application Data Grid Options**

| Option | Description |
|---|---|
| Name | Select the application name to display a summary of the application. Select to the right of the name to edit or delete the application. |
| Summary | Criticality of the health, risk, and efficiency of any application. Click an application with a red, orange, or yellow criticality to see more details about potential problems with objects in the application. |

### User Scenario: Adding an Application

As the system administrator of an online training system, you must monitor components in the Web, application, and database tiers of your environment that can affect the performance of the system. You build an application that groups

related objects together in each tier. If a problem occurs with one of the objects, it is reflected in the application display and you can open a summary to investigate the source of the problem further.

In your application, you add the DB-related objects that store data for the training system in a tier, Web-related objects that run the user interface in a tier, and application-related objects that process the data for the training system in a tier. The network tier might not be needed. Use this model to develop your application.

1. From the left menu, click **Operations › Applications**, then click **Applications** in the left pane.
2. Click **ADD**.
3. Click **Basic n-tier Web App** and click **OK**.
   The Application Management page that appears has two rows. Select objects from the bottom row to populate the tiers in the top row.
4. Type a meaningful name such as `Online Training Application` in the Application text box.
5. For each of the Web, application and database tiers listed, add the objects to the Tier Objects section.
   a) Select a tier name. This is the tier that you populate.
   b) To the left of the object row, select object tags to filter for objects that have that tag value. Click the tag name once to select the tag from the list and click the tag name again to deselect the tag from the list. If you select multiple tags, objects displayed depend on the values that you select.

      You can also search for the object by name.

   c) To the right of the object row, select the objects to add to the tier.
   d) Drag the objects to the Tier Objects section.
6. Click Save to save the application.

The new application appears in the list of applications on the Environment Overview Applications page. If any of the components in any of the tiers develops a problem, the application displays a yellow or red status.

To investigate the source of the problem, click the application name and see Evaluate Object Information Using Badge Alerts and Summary Tab .

To investigate the source of the problem, click the application name and evaluate the object summary information. See the *VMware Aria OperationsVMware Cloud Foundation Operations User Guide*.

### Add Application

When you add an application to an environment, you select from a list of predefined templates or create your own custom template, to group the objects to monitor in your application.

#### Where You Find Add Application

From the left menu, click **Operations › Applications**. From the Applications page, click **Add**.

#### Add Applications Options

Each predefined template provides you with a list of suggested tiers designed to help you group related objects that perform a specific task in your application. After you select an option, you can alter the selection and number of tiers on the Application Management page.

| Option | Description |
|---|---|
| Basic n-tier Web App | Use this template for any basic application. |
| Advanced n-tier Web App | Use this template for an application that monitors more physical devices, such as the devices that VMware Aria OperationsVMware Cloud Foundation Operations discovers when you add a network-related Management Pack or Management Packs. |
| Legacy non-Web App | Use this template for an application that has no Web-related objects. |
| Network | Use this template for an application that has only network-related objects. |
| Custom | Select this option to build your own application topology. |

### Application Management Dialog Box

You use Application Management to select the objects for your application. The objects you select are grouped in tiers and help you to track the health of your application.

#### Where You Find Application Management

From the left menu, click **Operations › Applications**. On the **Applications** tab, click **Add**. After you select an application template, click OK.

#### Application Management Options

At the top of the screen, enter a new application name or use the default name from the Add Application page. The application name must be unique.

Below the name, the page is divided into the tier row and the objects row. On each row, selections in the pane on the left filter the selections in the pane on the right.

The tier row is where you select the tiers to populate with objects to monitor for the application.

**Table 175: Tier Row**

| Option | Description |
|---|---|
| Tiers pane | Select the tier where you want to place your objects. You can add or delete tiers to fit your application. |
| Tier Objects pane | Add or remove objects that serve a common function and to monitor. For example, to monitor all the virtual machines that are database servers for the application, put them in the database tier. |

The object row is where you select objects to add to the tiers.

**Table 176: Object Row**

| Option | Description |
|---|---|
| Object Tags pane | Expand a tag to see a group of objects with that tag value. For example, if Adapter Types is an object tag, the tag values include vCenter Adapter, and an object is an adapter instance. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Objects are not displayed. The tag filters the object pane. To select a tag value, click once. To deselect a tag value, click twice. Tag values remain selected until they are deselected. |
| Objects pane | Drag an object with the object tag value to add to the Tier Objects pane. To find an object, search by name. Each object listed includes identifier information to help distinguish between objects of similar names. **Add All Objects To Parent** adds all the objects to a tier. |

## Dashboards in VMware Aria OperationsVMware Cloud Foundation Operations

Configuring Dashboards
Dashboards present a visual overview of the performance and state of objects in your virtual infrastructure. You use dashboards to determine the nature and timeframe of existing and potential issues with your environment. You create dashboards by adding widgets to a dashboard and configuring them.

VMware Aria OperationsVMware Cloud Foundation Operations collects performance data from monitored software and hardware resources in your enterprise and provides predictive analysis and real-time information about problems. The data and analysis are presented through alerts, in configurable dashboards, on predefined pages, and in several predefined dashboards.

- You can start with several predefined dashboards in VMware Aria OperationsVMware Cloud Foundation Operations.
- You can create extra ones that meet your specific needs using widgets, views, badges, and filters to change the focus of the information.
- You can clone and edit the predefined dashboards or start from scratch.
- To display data that shows dependencies, you can add widget interactions in dashboards.
- You can provide role-based access to various dashboards for better collaboration in teams.

**Table 177: Features**

| Features | Description |
|---|---|
| Manage | You can also manage dashboards by clicking **Operations › Dashboards**. From the Dashboards panel, click **Manage**. |
| Create | Use this option to create a dashboard. See Create and Configure Dashboards. |
| Search | You can search for a dashboard across the **Favorites**, **Recents**, and **All** folders in the **Dashboards** panel. |
| Favorites | You can mark a dashboard as a favorite using the **Favorite** icon at the top of each dashboard. All the dashboards that you have marked as a favorite, are listed under the Favorites folder in the **Dashboards** panel. |
| Recents | The dashboards are listed in the order in which you select them, with the most recent dashboard that you selected, appearing at the top. Up to ten dashboards can be displayed as Recent dashboards.<br>If you do not pin the dashboard and log out of the user interface, on logging back in, the dashboard is removed from the Recents folder. |
| Shared | If you have shared the dashboard, the shared icon is displayed against the dashboard name. |

*Table continued on next page*

*Continued from previous page*

| Features | Description |
|---|---|
| All | Lists the dashboard folders and the dashboards that are activated. You can use this menu for quick navigation through your dashboards. When you navigate to a dashboard using the **Operations › Dashboards** option, the dashboards are listed in the **Dashboards** panel under **All**. You can also search for dashboards using keywords and letters. |
| Actions | Available dashboard actions, such as edit, delete, Set as Dashboards Home, and Add to Product Home. These actions are applied directly to the dashboard that you are on.<br><br>**Set as Dashboards Home**: Adds the dashboard to the **Dashboard panel > Favorites** list. To remove the dashboard from the **Favorites** list, select **Actions › Unset as Dashboards Home**.<br><br>**Add to Product Home**: Adds the dashboard as a tab in the **Home** page. You can also reorder the tabs in the **Home** page using drag and drop. You can also set this option by clicking the **Home** icon at the top right side of the dashboard. To remove the dashboard from the Home page, select **Actions › Remove from Product Home**.<br><br>**NOTE**<br>You can add up to 5 dashboards to the **Home** page. |
| Dashboard Time | The dashboard time panel is activated by default on all predefined and user-created dashboards. Using this option, you can select a time for the widgets in the dashboard. The default time is 6 hours. The pre-defined time/day options in the panel are 1 hour, 6 hours, 24 hours, or 7 days. You can also set a customized time option.<br>To activate widgets to use the dashboard time, select ***Date Controls/Time Range*** **› Dashboard Time** from the widget toolbar. Some widgets have **Dashboard Time** as the default option. For example, **Metric Chart**, **View**, **Rolling View**, **Sparkline**, **Health Chart**, and **Mashup Chart** widgets.<br><br>Dashboard time persists if:<br>• You activate a widget in a dashboard to use the dashboard time and then log out and log back in, or<br>• You activate a widget in a dashboard to use the dashboard time, and you export and then import the dashboard into another instance of VMware Aria OperationsVMware Cloud Foundation Operations. |

## Accessing Predefined Dashboards

You can access some of the useful, predefined dashboards from the **Dashboards** home page.

To access these dashboards, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Home**.

The dashboards are categorized as follows: Availability, Configuration, Inventory, Performance, Capacity, and Cost. To easily access some of the useful, predefined dashboards under these categories, click on the drop-down button against the selected category and click on the specific dashboard.

## Types of Dashboards

You can use the predefined dashboards or create your own custom dashboard in VMware Aria OperationsVMware Cloud Foundation Operations.

See predefined dashboards for more information.

### Custom Dashboards

You can create dashboards that meet your environment needs in VMware Aria OperationsVMware Cloud Foundation Operations.

For information about creating a dashboard, see Create and Configure Dashboards.

## Create and Configure Dashboards

To view the status of all objects in VMware Aria OperationsVMware Cloud Foundation Operations, create a dashboard by adding widgets or views. You can create and modify dashboards and configure them to meet your environment needs.

1. From the left menu, click **Operations › Dashboards**.
2. From the **Dashboards** panel in the centre, click **Create**.
3. Complete the following steps to:
   a) Enter a name for the dashboard.

      Dashboard Name

   b) Add widgets or views to the dashboard.

      Widget or View List Details

   c) Configure widget interactions.

      Widget and View Interactions Details

   d) Create dashboard navigation.

      Dashboard Navigation Details
4. Click **Save**.
5. Click **Actions › Edit Dashboard** to modify the dashboard.

### Dashboard Name

The name and visualization of the dashboard as it appears on the VMware Aria OperationsVMware Cloud Foundation Operations Home page.

#### Where You Add a Name in a Dashboard

To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard. Enter a name in the **New Dashboard** field.

To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit Dashboard**.

If you use a forward slash while entering a name, the forward slash acts as a group divider and creates a folder with the specified name in the dashboards list if the name does not exist. For example, if you name a dashboard `clusters/hosts`, the dashboard is named `hosts` under the group `clusters`.

**Widget or View List Details**

VMware Aria OperationsVMware Cloud Foundation Operations provides a list of widgets or views that you can add to your dashboard to monitor specific metrics and properties of objects in your environment.

**Where You Add Widgets or Views to a Dashboard**

To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard.

To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit Dashboard**.

**How to Add Widgets or Views to a Dashboard**

In the widgets list panel, you see a list of predefined VMware Aria OperationsVMware Cloud Foundation Operations widgets or views that are most commonly used. Click **Show More** to view all the widgets and views. To add a widget or view to the dashboard workspace in the upper panel, you can:

- Drag the widget or view to the dashboard workspace in the upper panel, or
- Click on a widget or view from the widgets list panel, view a brief description of the widget or view from the pop up box and then click **Add to Dashboard**, or
- Double-click on the widget or view.

To locate a widget or view, you can enter the name or part of the name of a widget or view in the **Filter › Name** option. For example, when you enter `top`, the list is filtered to display the Top Alerts, Top-N, and Topology Graph widgets. You can then select the widget you require. You can also filter by either widgets or views by selecting **Filter › Show** to add a widget or view to the dashboard. Drag the widget or view to the dashboard workspace in the upper panel.

> **NOTE**
> The following widgets are deprecated. Usage of these widgets is discouraged as they will be phased out in future releases. Deprecated widgets are marked with a yellow triangle.
> - Current Policy
> - Weather Map
> - Anomalies
> - DRS Cluster Settings
> - Efficiency
> - Environment Status
> - Risk
> - Environment
> - Container Overview
> - Faults

Most widgets or views must be configured individually to display information. For more information about how to configure each widget, see Widgets in .

### How to Arrange Widgets or Views in a Dashboard

You can modify your dashboard layout to suit your needs. By default, the first widgets or views that you add are automatically arranged horizontally wherever you place them.

- To position a widget or a view, drag the widget or view to the desired location in the layout. Other widgets and views automatically rearrange to make room.
- To resize a widget or a view, drag the bottom-right corner of the widget or the view.
- To maximize or minimize a widget or a view, use the maximize and minimize options in the top-right corner.

## Widget and View Interactions Details

You can connect widgets and views so that the information they show depends on each other.

### Where You Create Widget and View Interactions

To create interactions for widgets or views in a dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard. From the toolbar, click **Show Interactions**.

To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit Dashboard**.

Internal interactions are interactions defined between widgets in a single dashboard. External interactions are interactions defined between a widget in one dashboard and widget/s in another dashboard.

### How to Create and Remove Widget Interactions

The list of available interactions depends on the widgets or views in the dashboard. Widgets and views can provide, receive, and can both provide and receive interactions at the same time.

To create interactions, click **Show Interactions**. Click a provider plug and drag to the receiver. You can also apply interactions from receiver to provider plugs. For more information about how interactions work, see Widget Interactions.

To remove interactions, click on the interaction line and select **Remove Interaction**. You can also click the provider plug and select **Remove Interaction › <widget name>**.

> **NOTE**
> You can create up to 25 widget interactions in a dashboard.

### How to Easily View Widget Interactions

To view widget interactions from a dashboard, click on an object from a provider widget that has widget interactions defined. A window appears with the external and internal widget interaction details. Clicking on an external interaction takes you to the external dashboard. Click on the internal interaction to view the details in the receiver widget.

## Dashboard Navigation Details

You can apply sections or context from one dashboard to another. You can connect widgets and views to widgets and views in the same dashboard or to other dashboards to investigate problems or better analyze the provided information.

### Where You Add Another Dashboard

To create dashboard navigation to a dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create** to add a dashboard. In the dashboard workspace, click **Show Interactions**. From the **Select Another Dashboard** drop-down menu, select the dashboard to which you want to navigate.

To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit Dashboard**.

### How Dashboard Navigation Works

You can create dashboard navigation only for provider widgets and views. The provider widget or view sends information to the destination widget or view. When you create dashboard navigation, the destination widgets or views are filtered based on the information type they can receive.

### How to Add Dashboard Navigation to a Dashboard

The list of available dashboards for navigation depends on the available dashboards and the widgets and views in the current dashboard. To add navigation, you can drag from a sender widget interaction plug to a receiver widget interaction plug. You can select more than one applicable widget or view.

> **NOTE**
> If a dashboard is unavailable for selection, it is unavailable for dashboard navigation.

The Dashboard Navigation icon ( ) appears in the top menu of each widget or view when a dashboard navigation is available.

After you have set widget interaction in the provider dashboard, the widget and menu bar are highlighted and two arrows appear in the top-left corner of the widget. After you have set widget interaction, clicking the object in the provider widget takes you to the receiver widget of the navigated dashboard.

## Manage Dashboards

You can select dashboards individually or as a group and perform several actions.

To manage your dashboards, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Manage**. Use the options from the horizontal ellipsis next to the **Add** option.

All the dashboards are listed on this page. You can filter the dashboards based on the name of the dashboard, the dashboard folder, activated dashboards, shared dashboards, or the dashboard owner. You can click **Add** to create a dashboard. For information about creating a dashboard, see Create and Configure Dashboards.

You can select a dashboard from the list, click the vertical ellipsis against each dashboard, and select the various options such as edit, delete, clone, and deactivate a dashboard. You can also change ownership of dashboards and export the dashboard. By default, the list of dashboards is sorted by name and all the columns can be sorted.

> **NOTE**
> A wrench icon appears when the data in an imported dashboard depends on the existence of one or more adapters that are currently not present. The wrench icon disappears if the required data in an imported dashboard appears in VMware Aria OperationsVMware Cloud Foundation Operations after configuration. Imported dashboards regardless of used data, remain stuck and include a wrench icon if the dashboard that is stuck (with the wrench icon), already exists.

**Datagrid Options**

| Column Names | Description |
|---|---|
| Name | Displays the name of the dashboard. |
| Folder | Lists the folder to which each dashboard belongs. |
| Description | Displays the description of the dashboard. |

*Table continued on next page*

*Continued from previous page*

| Column Names | Description |
|---|---|
| Activated | Displays whether the dashboard is activated or not. You can also activate and deactivate the dashboard. |
| URL | Displays whether the dashboard is shared externally. For dashboards that have been shared, click to view the shared links. |
| Shared | Displays whether the dashboard is shared internally. Click the icon to view and edit the groups to which the dashboard has been shared from the **Group Sharing** dialog box.<br>From the **Group Sharing** dialog box, to share dashboard edit privileges with the group, you can click the **Allow Editing** check box or click the pencil icon if a dashboard is shared with a group.<br>**NOTE**<br>Users who are part of the user group to whom the dashboard has been shared, will see an unlocked or locked icon in the **Shared** column. The icon is unlocked when the user with whom the dashboard is shared can edit it, and the icon is locked when the dashboard cannot be edited. |
| Editable | Displays if the dashboard is editable. |
| Owner | Displays the owner of the dashboard. |
| Report Usage | Displays the number of reports where the dashboard is used. Click the number in the column to view the name of the report/s. Click the report name to navigate to the report template in edit mode. |
| Last Modified | Displays the date the dashboard was last modified. |
| Modified By | Displays the user who last modified the dashboard. |

You can select more than one dashboard and perform a set of options by clicking the horizontal ellipsis next to the **Add** option.

**Table 178: Dashboards Options**

| Option | Description | Usage |
|---|---|---|
| Delete | Deletes a dashboard. | |
| Activate | Activates a dashboard that was previously deactivated. | |
| Deactivate | Deactivates a dashboard. | |
| Change Ownership | Assigns a new owner to the dashboard. | After you assign a dashboard to a new owner, the dashboard is no longer displayed as one of your dashboards.<br>When you transfer a dashboard that was previously shared with user groups, information |

*Table continued on next page*

*Continued from previous page*

| Option | Description | Usage |
|---|---|---|
| | | about the shared user groups and group hierarchy is retained. |
| Export | When you export a dashboard, VMware Aria OperationsVMware Cloud Foundation Operations creates a dashboard file in JSON format. | You can export a dashboard from one VMware Aria OperationsVMware Cloud Foundation Operations instance and import it to another.<br>To export a dashboard, select the dashboard that you want to export, and click **Export** from the horizontal ellipsis. |
| Import | A PAK or JSON file that contains dashboard information from VMware Aria OperationsVMware Cloud Foundation Operations. | You can import a dashboard that was exported from another VMware Aria OperationsVMware Cloud Foundation Operations instance.<br>To import a dashboard:<br>1. Click the **Import** option from horizontal ellipsis.<br>2. Click **Browse** and select a Dashboard ZIP, PAK, or JSON file to import.<br>3. Select if you want to **Overwrite** or **Rename** the file in case of a conflict.<br>4. Click **Import** to import the dashboard, and click **Done**. |
| Auto-rotate Dashboards | Changes the order of the dashboard tabs on VMware Aria OperationsVMware Cloud Foundation Operations home page. | You can configure VMware Aria OperationsVMware Cloud Foundation Operations to switch from one dashboard to another. For more information, see Auto-Rotate Dashboards. |
| Manage Summary Dashboards | Provides you with an overview of the state of the selected object, group, or application. | You can change the **Summary** tab with a dashboard to get information specific to your needs. For more information, see Manage Summary Dashboards |
| Manage Dashboard Folders | Groups dashboards in folders. | You can create dashboard folders to group the dashboards in a way that is meaningful to you. For more information, see Manage Dashboard Folders. |
| Manage Dashboard Sharing | Makes a dashboard available to other users or user groups. | You can share a dashboard or dashboard template with one or more user groups. For more information, see Share Dashboards with Users. |

The dashboard list depends on your access rights.

**Manage Summary Dashboards**

The **Summary** tab provides you with an overview of the state of the selected object, group, or application. You can change the **Summary** tab with a dashboard to get information specific to your needs.

**Where You Configure a Summary Tab Dashboard**

To manage the summary dashboards, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Manage**. Click the horizontal ellipsis next to the **Add** option and select **Manage Summary Dashboards**.

### How You Manage the Summary Dashboards

**Table 179: Manage Summary Dashboards Toolbar Options**

| Option | Description |
|---|---|
| Use Default | Click to use VMware Aria OperationsVMware Cloud Foundation Operations default **Summary** tab. |
| Assign a Dashboard | Click to view the Dashboard List dialog box that lists all the available dashboards. |
| Adapter Type | Adapter type for which you configure a summary dashboard. |
| Filter | Use a word search to limit the number of adapter types that appear in the list. |

To change the Summary tab for an object, select the object in the left panel, click the **Assign a Dashboard** icon. Select a dashboard for it from the All Dashboards dialog box and click **OK**. From the Manage Summary Dashboards dialog box click **Save**. You see the dashboard that you have associated to the object type when you navigate to the **Summary** tab of the object details page.

## Auto-Rotate Dashboards

You can change the order of the dashboard tabs on your home page. You can configure VMware Aria OperationsVMware Cloud Foundation Operations to switch from one dashboard to another. This feature is useful if you have several dashboards that show different aspects of your enterprise's performance and you want to look at each dashboard in turn.

### Where You Configure Auto-Rotation of a Dashboard

To reorder and configure a dashboard switch, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Manage**. Select **Auto-rotate Dashboards** from the horizontal ellipsis next to the **Add** option.

### How You Reorder the Dashboards

The list shows the dashboards as they are ordered. Drag the dashboards up and down to change their order on the home page.

### How You Configure an Automatic Dashboard Rotation

1. Double-click a dashboard from the list to configure.
2. From the Rotation drop-down menus, select **On**.
3. Select the time interval in seconds.
4. Select the dashboard to switch and click **Update**.
5. Click **Save** to save your changes.

On the home page, the current dashboard will switch to the dashboard that is defined after the specified time interval.

## Manage Dashboard Folders

You can create dashboard folders to group the dashboards in a way that is meaningful to you.

### Where You Manage Dashboard Folders

To manage the dashboard folders, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Manage**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Folders**.

### How You Manage the Dashboard Folders

**Table 180: Manage Dashboard Folders Options**

| Option | Description |
|---|---|
| Dashboards List | A list with all available dashboards. |
| Folders | A hierarchy tree with all the available group folders. |

To create a dashboard folder, click **New Folder** in the **Folders** pane and enter the name of the folder. If you want to create a folder under another folder, select a parent folder under which you want to create the child folder, then click **New Folder**. To add a dashboard, drag one from the dashboards list to the selected folder in the **Folders** pane.

You can delete folders and/or detach dashboards from a folder, by selecting one or more folders and dashboards from the **Folders** pane and by clicking **Actions › Delete**.

You can rename a folder by selecting a single folder from the **Folders** pane and by clicking **Actions › Rename**.

## Share Dashboards with Users

You can share a dashboard with one or more user groups. When you share a dashboard, it becomes available to all the users in the user group that you select. The dashboard appears the same to all the users who share it. If you edit a shared dashboard, the dashboard changes for all users. Other users can only view a shared dashboard. They cannot change it.

### Where You Share a Dashboard From

To share a dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Manage**. Click the horizontal ellipsis next to the **Add** option and click **Manage Dashboard Sharing**.

**Table 181: Dashboard Sharing Options**

| Option | Description |
|---|---|
| All Dashboards | Link to view all the available dashboards that you can share. The dashboards are displayed on the right side in the dashboards list. |
| User Groups | Lists the available user groups that you can share a dashboard with. The list includes the **Everyone** group. |
| Dashboard List | List of shared dashboards with the selected user group or all the available dashboards that you can share, if no user group is selected. |

### Manage Dashboard Sharing

To share a dashboard, navigate to the dashboard in the list of dashboards and drag it to the group to share it with, on the left.

To stop sharing a dashboard with a group, click that group on the left panel, navigate to the dashboard in the right panel, and click **Stop Sharing** above the list.

**Table 182: Datagrid Options**

| Column Name | Description |
|---|---|
| Name | Displays the name of the dashboard. |
| Folder | Displays the folder where the dashboard is located. |
| Description | Displays the description of the dashboard. |
| Activated | Displays whether the dashboard is activated for viewing. |
| Shared | Displays if the dashboard is shared with user groups. |
| Editable | Click the check box to share dashboard edit privileges with the specific group. |
| Owner | Displays the owner of the dashboard. |
| Last Modified | Displays the date the dashboard was last modified. |
| Modified By | Displays who modified the dashboard. |

## Dashboards Actions and Options

You can change the order of the dashboard tabs, configure VMware Aria OperationsVMware Cloud Foundation Operations to switch from one dashboard to another, create dashboard folders to group the dashboards in a way that is meaningful to you, share a dashboard or dashboard template with one or more user groups, and transfer selected dashboards to a new owner.

### Options for Sharing Dashboards

You can share predefined or custom dashboards using URLs, emails, and by copying the code to embed the dashboard into confluence or other internal official web pages. You can also assign and unassign a dashboard to specific user groups and export the dashboard configuration details. You can share predefined or custom dashboards by assigning and unassigning a dashboard to specific user groups and by exporting the dashboard configuration details.

When you use a non-authenticated shared URL, as a user you can open the dashboard in a new browser session. If you have already logged into VMware Aria OperationsVMware Cloud Foundation Operations in another session, you are redirected to this dashboard and the user authentication permissions apply. To ensure that the non-authenticated URL opens the intended dashboard, as a user you must log out from all existing user sessions.

The dashboard shared with the URL opens in a page where you can access all the widgets within the dashboard and you can interact with the given widgets at the same time. A non-authenticated dashboard however, does not allow you to browse to other areas of VMware Aria OperationsVMware Cloud Foundation Operations.

Dashboard sharing can only be applied to Groups with a VMware Aria Operations Standard Edition license.

#### Where You Can Access the Options to Share Dashboards

From the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click on an existing dashboard and then click the **Share Dashboard** icon in the top-right corner.

**Table 183: Options in the Share Dashboard Dialog Box**

| Option | Description |
|---|---|
| URL | Allows you to copy the tiny URL for the selected dashboard.<br>• Set the expiry period for the link to **1 day**, **1 week**, **1 month**, **3 Months**, or **Never Expire**. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Click **Copy Link** to copy the link to a new window from where you can view the dashboard.<br><br>**NOTE**<br>• As a user, if you open a shared link and you are logged into VMware Aria OperationsVMware Cloud Foundation Operations, you are navigated to your default dashboard, instead of viewing the shared one.<br>• As a user, if you log in to the same IP that was shared with you previously, you cannot access the page with the same browser.<br>• As a user, ensure that you have the following permission: **Dashboards > Dashboard Management > Share (Public)**.<br><br>You can stop sharing a dashboard you had previously shared. To stop sharing a dashboard, click the **Unshare Link** option and enter the URL of the dashboard that you want to stop sharing and click **Unshare**.<br>Authentication is not required to view the shared dashboard. |
| Email | Allows you to send an email with the URL details of the dashboard, to a specific person.<br>• Set the expiry period for the link to **1 day**, **1 week**, **1 month**, **3 months**, or **Never Expire**.<br>• Configure an SMTP instance. See Add a Standard Email Plug-In for Outbound Alerts.<br>• Enter an email address and click the **Send Email** button to send an email with the URL details of the dashboard.<br><br>Authentication is not required to view the shared dashboard. |
| Embed | Provides an embedded code for the dashboard. You can use this code to embed the dashboard in relevant confluence pages that your company executives routinely use and analyze.<br>• Set the expiry period for the link to **1 day**, **1 week**, **1 month**, **3 Months**, or **Never Expire**.<br><br>**NOTE**<br>• If you embed a dashboard in the **Text** widget, the widget does not display any data.<br>• When you open an HTML/confluence page with an embedded dashboard from the same browser that you have logged into VMware Aria OperationsVMware Cloud Foundation Operations, the dashboard does not load.<br><br>Authentication is not required to view the shared dashboard. |
| Groups | Allows you to assign and unassign a dashboard to specific user groups.<br>• Select the group to which you want to grant dashboard access from the drop-down menu. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • To share dashboard edit privileges with the group, select the **Allow Editing** check box if a dashboard is shared with a group.<br>• Click **Include**. You can include more than one group.<br>• From the label, select the cross mark to unassign the dashboard.<br><br>Log in to VMware Aria OperationsVMware Cloud Foundation Operations to view the shared dashboard. |
| Export | Allows you to export the dashboard configuration details.<br>Log in to VMware Aria OperationsVMware Cloud Foundation Operations to export/import a dashboard. |

**Table 184: Options in the Share Dashboard Dialog Box**

| Option | Description |
|---|---|
| Groups | Allows you to assign and unassign a dashboard to specific user groups.<br>• Select the group to which you want to grant dashboard access from the drop-down menu.<br>• To share dashboard edit privileges with the group, select the **Allow Editing** check box if a dashboard is shared with a group.<br>• Click **Include**. You can include more than one group.<br>• From the label, select the cross mark to unassign the dashboard.<br><br>Log in to VMware Aria OperationsVMware Cloud Foundation Operations to view the shared dashboard. |
| Export | Allows you to export the dashboard configuration details.<br>Log in to VMware Aria OperationsVMware Cloud Foundation Operations to export/import a dashboard. |

### Manage Widgets in Dashboards

You can replicate widgets multiple times in a dashboard by using the copy and paste functionality.

Navigate to the dashboard from which you want to copy widgets. Select **Actions › Edit Dashboard**. Select one or more widgets that you want to copy by clicking the title of the widget and then select **Actions › Copy Widget(s)**. Click **Actions › Paste Widget(s)** to paste one or more widgets in the same dashboard.

To paste one or more widgets into another dashboard, exit the edit screen of the dashboard by selecting **Cancel**. Navigate to the dashboard to which you want to paste one or more widgets and select **Actions › Edit Dashboard** and then **Actions › Paste Widget(s)**.

## Predefined Dashboards

VMware Aria OperationsVMware Cloud Foundation Operations includes a broad set of simple to use, but customizable dashboards to get you started with monitoring your VMware environment. The predefined dashboards address several key questions including how you can troubleshoot your VMs, the workload distribution of your hosts, clusters, and datastores, the capacity of your data center, and information about the VMs. You can also view log details.

Each set of dashboards is complemented with a series of out-of-the-box customizable alerts and reports to assist with your operational awareness. Alerts, reports, and dashboards, each have a purpose with minimal overlap. Several activities that are carried out using alerts should be carried out using dashboards. Reports should be kept to a minimum as they are not interactive and do not provide timely information.

The following table details how alerts, dashboards, and reports are complimentary.

| | Alerts | Dashboards | Reports |
|---|---|---|---|
| Nature | Reactive | Proactive | Passive.<br>For those with no access to<br>VMware Aria Operations<br>and VMware Aria Operations for Logs. |
| Suitability | Exception<br>(something went wrong) | Exception<br>Big Picture<br>Details Analysis | Big Picture<br>Exception (but not urgent)<br>No analysis as it is not interactive |
| Use Case | Troubleshooting (the start) | Monitoring<br>Troubleshooting (the actual) | FYI (optional)<br>Export for further analysis (spreadsheet) |
| Time & Urgency | Urgent (minutes) &<br>Important | Regular (daily, SOP) | Not urgent (monthly) & optional<br>No point in a daily report.<br>For daily reports, login for<br>interactivity |
| Access Requirement | Online. Desktop.<br>1280 * 1024 pixel | Online. Desktop.<br>280 * 1024 pixel | Offline or mobile. Small resolution.<br>Email. Laptop or tablet. |
| Scope/Area | Availability<br>Performance<br>Compliance<br>Configuration (?)<br>Capacity (less relevant,<br>unless it is an emergency) | Availability<br>Performance<br>Capacity<br>Compliance<br>Configuration<br>Inventory | Same as dashboards, but:<br>• without interactivity<br>• time bound (e.g. calendar month)<br>• No performance report, covered in<br>Capacity |
| Roles | Operations Team | Operations Team<br>Architect Team | IT Management (not hands-on)<br>Auditor (compliance) |

**Insight vs Alerts**

VMware Aria OperationsVMware Cloud Foundation Operations dashboards support a concept we call insight. Insight complements alerts but does not replace it. Alerts miss the larger picture and only see what is triggered. For one object that reaches the threshold, there might be many just beneath the threshold. The objects below the threshold are called insight.

Alerts might auto-close if the symptoms disappear. Managing alerts is not the same as minimizing alerts. Minimizing alerts is about preventing alerts.

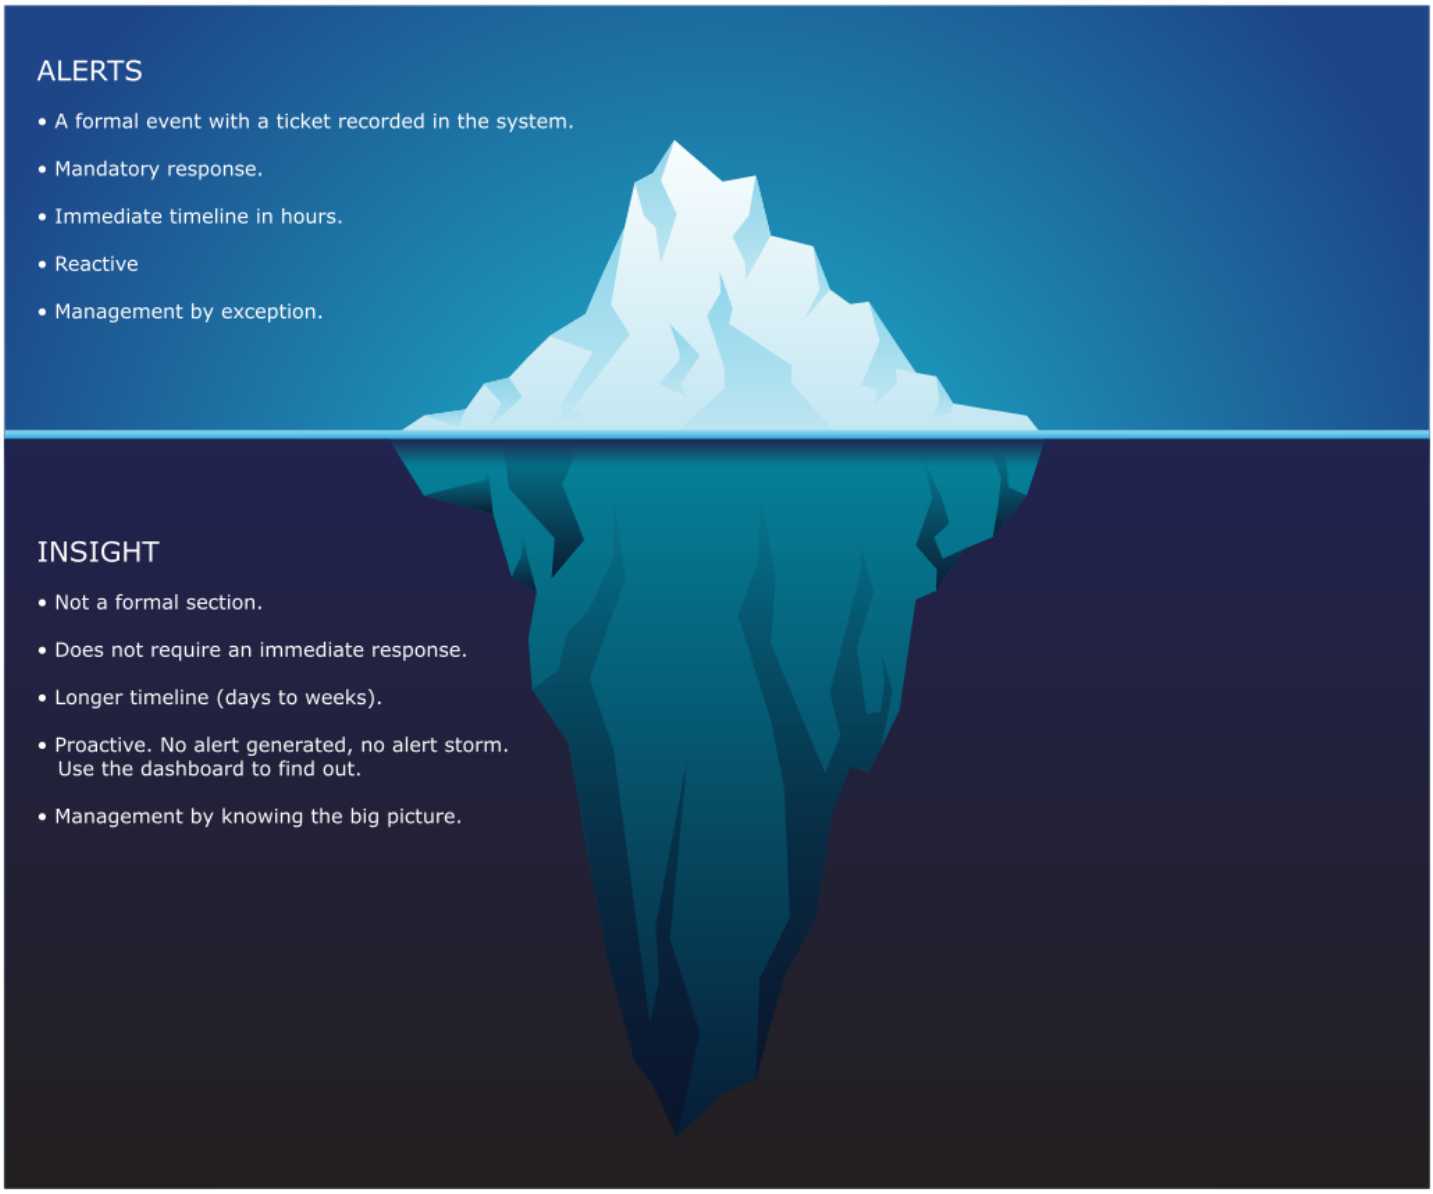


### Working with Predefined Dashboards

To access the predefined dashboards, from the left menu, click **Operations › Dashboards**.

You can customize dashboards and widgets if you have VMware Aria Operations Advanced edition or higher. Any customization you make is overwritten during upgrade and as a result, it is recommended that you back up your dashboards before an upgrade.

## Application Monitoring Dashboards

Use the **Linux OS discovered by Telegraf** dashboard and the **Windows OS discovered by Telegraf** dashboard to view the Linux and Windows OS objects discovered on all types of machines such as Virtual Machines, Azure Virtual Machine, EC2 Instance, and physical servers.

### Linux OS discovered by Telegraf Dashboard

Use the **Linux OS discovered by Telegraf** dashboard to view detailed information related to Linux OS objects. You can monitor CPU, disk, and network metric usage. You can also view object relationships.

**Table 185: Widgets in the LInux OS discovered by Telegraf Dashboard**

| Widget | Description |
|---|---|
| Linux Operating Systems | Displays a list Linux OS objects that are discovered on all types of machines. Click on the Linux OS object for which you want to view details. The other widgets are populated with relevant details. |
| Process | Displays Linux Process metrics for the selected Linux OS object. |
| CPU | Displays CPU metrics for the selected Linux OS object. |
| Memory | Displays Memory metrics for the selected Linux OS object. |
| Memory Swap File | Displays Memory Swap File metrics for the selected OS object. |
| CPU: Individual CPU Performance | Displays all CPU usage metrics for the selected Linux OS object. |
| Disk: Individual Partition Capacity | Displays the instance name, total capacity, free capacity, and percentage of capacity that is used. |
| Disk: Individual Disk Performance | Displays the Individual disk performance metrics for the selected Linux OS object. |
| Related Objects | Displays a detailed object relationship of the Linux OS objects. |

### Windows OS discovered by Telegraf Dashboard

Use the **Windows OS discovered by Telegraf** dashboard to view detailed information related to Windows OS objects. You can monitor CPU, disk, and network metric usage. You can also view object relationships.

**Table 186: Widgets in the Windows OS discovered by Telegraf Dashboard**

| Widget | Description |
|---|---|
| Windows Operating Systems | Displays a list Windows OS objects that are discovered on all types of machines. Click on the Windows OS object for which you want to view details. The other widgets are populated with relevant details. |
| Associated VM | Displays the VM version, CPU, memory, and disk space details. |
| System | Displays System metrics for the selected Windows OS object. |
| CPU | Displays CPU metrics for the selected Windows OS object. |
| Memory | Displays Memory metrics for the selected Windows OS object. |

*Table continued on next page*

*Continued from previous page*

| Widget | Description |
|---|---|
| Windows: Memory Rate | Displays Memory Rate metrics for the selected Windows OS object. |
| CPU: Individual CPU Usage | Displays all CPU usage metrics for the selected Windows OS object. |
| Disk: Individual Partition Capacity | Displays the instance name, total capacity, free capacity and percentage of capacity that is used. |
| Network: Individual Adapter Usage | Displays the individual network adapter usage metrics for the selected Windows OS object. |
| Related Objects | Displays a detailed object relationship of the Windows OS objects. |

## Availability Dashboards

Availability covers the uptime of the object now and the uptime trend over time. The availability of hybrid clouds should be tracked at both the provider and consumer layers to understand the availability of the environment. These dashboards show the current uptime and the uptime percentage over the past month.

### VM Availability Dashboard

Use the **VM Availability** dashboard to calculate the availability of the Guest OS. The availability of the Guest OS is calculated because the Guest OS might not be running even when the VM is powered on. There are two layers of Availability, that is, the Consumer layer and the Provider layer. This dashboard covers the Consumer layer. You can view VMs in the selected data center, uptime trend for a selected cluster, and so on.

#### Design Considerations

The **VM Availability** dashboard helps you check the availability (uptime in percentage) of VMs, as availability is typically part of the services provided by the IaaS provider.

This dashboard does not check the application uptime because it is possible that the application such as, a database, or a web server, is down while the underlying Windows or Linux is up. Generally, the service provided by the IaaS team is only for Windows or Linux. For information on the application, use the network ping or application-specific agent such as application monitoring.

#### How to Use the Dashboard

- In the **Datacenters** widget, click any data center from the list.
  - To view the overall information, click the **vSphere World** object.
  - The other widgets are automatically updated once you click any data center.
  - Create a filter that reflects your class of service for this widget. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads and you can focus on the important VMs. You can achieve this by creating a VMware Aria OperationsVMware Cloud Foundation Operations custom group for each class of service.
- The **VMs by Uptime in the last 30 days** widget displays the average uptime of VMs grouped by their availability. The bucket distribution helps you cater to a wide array of environments. If you are monitoring only production VMs where the uptime is expected to be near 100% all the time, edit the bucket to meet your operational needs.
  - The VMs in the **Selected Datacenter** widget display all the VMs that are currently deployed to the data center. The average uptime is displayed for the last month. For a production VM, expect this number to be 100% or closer to 100%.

**NOTE**
The Services column will be blank unless Service Discovery is activated and the services/processes are discovered on a specific virtual machine.

  - The VMs column includes all VMs including the powered off VMs.
- Click any VM in the **VMs by Uptime in the last 30 days** widget to view the details in the **VM in the Selected VM Powered On Status**, **Selected VM Uptime Trend**, and **Selected Cluster Uptime Trend** widgets.
  - The **Selected VM Uptime Trend** widget displays the selected VM's Guest Tool Uptime (%) across the last 30 days.
- The **Guest OS: Services** widget displays the service state over time and the process or services running inside the Guest OS. If Guest OS services or processes are discovered inside a VM, their availability is analyzed. This requires the Service Discovery.
- The **ESXi Host(s) where the VM has run** widget displays the historical migration of the VM. This can be useful in determining the cause of a VM downtime.

**Points to Note**
- The metric only tracks the availability of VMware Tools and not the entire Guest OS. If VMware Tools is not up, it assumes the Guest OS to be down. You can check that this is not a false negative by adding a few line charts that display the evidence of activity. A good counter is IO counters such as Disk IOPS, Disk Throughput, and Network Transmit Throughput, because IO requires CPU processing. CPU usage is not a reliable counter as the work by VMkernel on the VM is charged to the CPU counters.
- VMware Aria OperationsVMware Cloud Foundation Operations exhibits a new ping adapter. This allows you to enhance the accuracy of the uptime measurement by creating a super metric that adds the ping information or by checking the process using an agent, such as application monitoring.
- Add a property widget that lists the selected VM properties to give you more context about the VM. In a large environment, the VM name alone might not provide enough context.

### vSphere Availability Dashboard

There are two layers of Availability, that is, the Consumer layer and the Provider layer. The **vSphere Availability** dashboard covers the Provider layer. This dashboard includes a cluster and not an ESXi host because the cluster is operationally a single compute provider. This dashboard considers the N+1 design, where the cluster can withstand one host failure. Logically, a cluster with fewer hosts has a higher risk.

**Design Considerations**

The **vSphere Availability** dashboard helps you analyze and report the uptime, as availability is typically part of the official business SLA. It is also often required in the monthly operational summary report.

This dashboard is not designed for live monitoring of the uptime. A NOC style of dashboard is better suited for those use cases. VMware Tools such as VMware Cloud Foundation Operations for logs must be leveraged as the fault is typically preceded with soft errors.

**How to Use the Dashboard**
- The **Clusters** widget lists all the clusters in the environment. It is sorted by the lowest uptime so that the cluster with the lowest uptime in the last one month is displayed.

  - The **Running Hosts** column is color-coded as logically a smaller cluster has a higher risk. A single host failure results in a relatively higher capacity degradation.
  - The **vSAN?** column is hyper-converged, which means both the compute and the storage part is considered.
  - The **Admission Control Policy** column is based on the Cluster Configuration \ DAS Configuration \ Active property. The mapping between the code to name is:
    - -1 : Disabled
    - 0 : Cluster Resource Percentage
    - 1 : Slot Policy (Powered-on VMs)
    - 2 : Dedicated Failover Hosts
  - In a large environment, creating a filter for the list of clusters can make it more manageable. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, you can easily view your gold clusters.
- Click any cluster from the **Clusters** widget.
  - The cluster uptime is automatically plotted in the **Selected Cluster Uptime Trend** widget. It uses 99%, 99.%, and 99.99% as the threshold for red, orange, and yellow colors respectively.
  - The ESXi host details in **ESXi in the Selected Cluster** widget are automatically updated. For more context, you can add a property widget that lists the selected ESXi host properties.
  - In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.
- The **Datastores not available** widget lists only the datastores with powered off status. This covers both local and shared datastores. To add context, consider adding an extra column such as the data center where it resides, and the datastore types such as NFS and VMFS.
- The **Port Group Availability** widget lists port groups that currently have an uptime of less than 100%. To add context, consider adding an extra column such as the data center where it resides, number of used ports, and the maximum number of ports.
- For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.
- In a large environment, you can create a filter for this dashboard. Group by the class of services such as gold, silver, and bronze and default the selection to Gold. In this way, the monitoring is not cluttered with less critical workloads.
- In the **ESXi in the Selected Cluster** widget, the **Connected to vCenter** and **Maintenance State** columns are not the average values, as both are string. However, they display the last state in the selected period. This allows you to go back to a specific point in time and view availability at that point.

**Points to Note**

- You can add vCenter Server and NSX components availability. This requires the VMware SDDC Health Monitoring Solution.

### Ping Overview Dashboard

Use the Ping Overview dashboard to configure the ping functionality and verify the availability of end points that exist in your virtual environment. The ping functionality is configured at the adapter instance for IP addresses, group of IP addresses, and FQDN. You can view ping adapter details like, latency distribution and packet loss distribution in this dashboard.

**Customizations Available for Your Use**

For more context, you can add a property widget that lists the selected object properties. Multiple tables can drive the same property widget, but the object type must be the same.

**NOTE**
The FQDN names are checked for validity, the FQDN validation relies on RFC1034 and RFC1123, and only top level domains of the internet are validated. The `.local` domain is not supported as it does not fall into the list of top-level domains in the Domain Name System (DNS) of the Internet.

**Widget Information**

- Latency distribution - You can use this widget to see the objects that are experiencing high latency.
- Packet Loss Distribution – You can use this widget to see the objects that are experiencing high packet loss.
- Ping Targets – You can use this widget to view the list of ping targets grouped by their FQDN. Latency and packet loss information is also displayed for the ping objects.
- Breakdown by Source Initiator – You can use this widget to view the List of ping statistics by the source (ping initiator). You can ping the target from multiple locations, to determine if the issue is network-related or server-related.

### vSAN Health Dashboard

Use the vSAN Health dashboard to monitor the health of your vSAN cluster.

You can use the vSAN Health dashboard to monitor the status of the cluster components, diagnose issues, and troubleshoot problems. The dashboard displays the critical and important alerts that need immediate attention.

**How to Use the Dashboard**

The **vSAN Clusters** widget lists all the clusters in the environment with issues if the cluster health score is below 60 or has any critical alerts.

It shows the cluster names, ESXi hosts, resync, and color-coded health and risk of the cluster. vSAN resync is a utilization metric, but its presence can impact performance. If the table is all green, no need to analyze it further.

You can focus on issues that need immediate attention and resolution for your clusters to be up and running. This dashboard displays the alerts triggered in your vSAN cluster and provides insights into the cause of the alerts.

The **Red Alert** column displays the critical alerts in the cluster that need immediate attention. If there are red alerts, click the cluster. The **Alerts** tab displays all the alert lists. Click a particular alert and check the alert details, related alerts, and potential evidence related to that alert. Under **Alert Details**, you can check the recommended remediation to troubleshoot the specific issue. Click the Knowledge Base article for more details.

Click any cluster to view the properties of the selected cluster in the **Property of Selected Cluster** widget. These are the properties that helps you understand the cluster health issues. If something is not there, check the alerts in the alert lists.

Once you select any particular cluster, the alerts in that cluster are listed in the **Alert Lists** widget. Click an alert to view the alert details, related alerts, and potential evidence. It also displays the alert symptoms.

## Capacity Dashboards

Capacity quantifies the resources used, resources remaining, and opportunities to reclaim unused resources. Projections of the demand provide a proactive view of capacity. The **Capacity Dashboards** display capacity in terms of time remaining before capacity is projected to run out, the amount of capacity remaining, the number of VMs that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

Capacity management is about balancing demand and supply. It is about meeting demand with the lowest possible cost.

For IaaS or DaaS, capacity management begins before the hardware is deployed. It begins with a business plan that defines what class of service will be provided. Each class of service, for example, gold, silver, bronze is differentiated by

the quality of service and covers the availability, for example, 99.99% uptime for Gold, 99.95% uptime for Silver. It also covers performance, for example, 10 ms disk latency for Gold, 20 ms disk latency for Silver, and security or compliance.

The quality incurs cost and in turn drives price. Gold VM is higher per vCPU and per GB of RAM because it has a higher quality of service. A proper pricing model must be planned. If you want your customers to rightsize in advance, then a 64 vCPU VM has to be more than 64x the price of 1 vCPU VM. If the pricing model is a simple straight line, there is no incentive to go small and no penalty if it is over provisioned. In this case, you end up forcing rightsizing in production, which is a costly and time consuming process.

Demand is more than the active load that is consuming your capacity. Since capacity based on utilization is incomplete by itself, the principles displayed in the following figure are considered.


- Latent demand. Many critical VMs are protected with Disaster Recovery. During a Disaster Recovery drill or an actual disaster, this load is consumed.
- Potential demand. Many newly provisioned VMs take time to reach their expected demand. It takes time for the database to reach the full size, the user base to reach the target, and the functionalities to complete. When this is achieved, it results in the increase in demand.
- Unmet demand occurs when the VM or Kubernetes Pod is undersized. The load is running nearly 100% most of the time.
- Excessive demand can wreak havoc in a shared environment. A group of highly demanding VM can collectively impact overall performance of the cluster or datastore.

**Cluster Capacity Dashboard**

The **Cluster Capacity** dashboard includes the ESXi host and resource pools as they impact cluster capacity.

**Design Considerations**

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

**How to Use the Dashboard**

The **Cluster Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

**Overall Analysis**

- The three bar charts which are **Clusters by Capacity Remaining**, **Clusters by Time Remaining**, **Clusters by VM Remaining**, summarize the overall situation. The first two charts can be used together to identify when you need to add capacity to address growth. Time remaining uses historical growth in a cluster to forecast when more capacity is needed. This allows you to operate more efficiently by making sure you have enough capacity currently and proactively plan for adding capacity. The third bar chart which is **Clusters by VM Remaining**, provides complete contexts, as different clusters can have different VM sizes.
  For a large environment, a heat map is helpful. The three heat maps are Time Remaining, Capacity Remaining, and VM Remaining. If your cluster sizes are not standardized, create another heat map, and use the number of ESXi hosts to show the size difference.

**Cluster Analysis**

- The **Clusters Capacity** widget provides a table with details. The number of ESXi hosts are color coded as smaller clusters have a relatively higher overhead. Select a cluster from the table to view the capacity details that are automatically displayed.

  **Performance**

  Ensure that the performance of the cluster meets your SLAs.

  **Utilization**

  The next two charts are Memory Workload (%) and CPU Workload (%), that show values relative to your usable capacity. Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average, so you can focus on the overall trend. For memory, the focus is on consumed memory and not active memory.

  **Allocation**

  You can view the trend of the three which are CPU, disk, and memory components together on the **Overcommit Ratio** chart. In general, your CPU overcommit should be the highest, followed by the disk (because of thin provision). Memory overcommit tends to be near one due to its nature as cache.

  Use the line chart in the **Allocation** widget, to see the trend. The data is averaged hourly.

  In the **VM Count** widget, the trend line of the number of VMs over time is important to spot if there are many newly provisioned VM. If you see that the VMs are increasing but demand remains low, it indicates a sign of potential demand in the future.

  **Reservation**

  Reservation can impact the efficiency of your cluster. Your cluster could be low on capacity because of real workload or just reservation. If your cluster size varies, complement the reservation number by showing a relative value. Once you have a standardized number, you can visualize them on a heat map.

- **ESXi Analysis**
  Good cluster capacity does not indicate that there is no issue at the ESXi level. Unbalance is a common problem, especially in large clusters and stretched clusters.

  The **ESXi Hosts in a Cluster** table displays all the member ESXi hosts. You can see the unbalance clearly, thanks to the color code. The color code reflects the unbalance.

  The **99th percentile Performance** column takes the 99th percentile value of the ESXi Performance (%) metric.

  Select an ESXI host to view the details. Both the **CPU Workload (%)** and the **Memory Workload (%)** trend line charts display if there is a steady demand, cyclical demand, rising demand, or declining demand. The trend is as important as the present value. View trends over a longer time. Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average. The focus is on memory consumed and not memory active. Memory consumed includes the total memory consumed, so it includes the memory consumed by

VMkernel. Both total and usable utilization in terms of memory and CPU are displayed and provides the absolute amount of capacity.

- **VM Analysis**
  Use the **VMs in the selected Cluster or Host** table to analyze the cause of the low capacity remaining and which VMs are impacting the infrastructure resources, such as, CPU, memory, and disk space. The table lists either the VMs in the cluster or host. When you select one of the VMs, additional relevant information is displayed.

  If there are many large VMs running low on capacity you can stop provisioning until you upsize the existing VMs first.

### Datastore Capacity Dashboard

The **Datastore Capacity** dashboard complements out of the box capacity pages and dashboards. It focuses on storage, provides an overall picture, and highlights the datastores that need attention.

#### Design Considerations

See Capacity Dashboards for common design consideration among all the dashboards for capacity management.

#### How to Use the Dashboard

The **Datastore Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

- **Overall Analysis**
  The summary banner answer basic questions such as number of datastores, the capacity, number of VMs, and the running VMs.

  The distribution charts are **Shared Datastores by Capacity Remaining** and **Shared Datastores by Time Remaining**.

  There are three heat maps, the primary being the **Remaining Capacity** heat map. The two other heat maps cover used capacity. One of them is designed for an environment that uses datastore clusters. Each box represents a datastore. If you have many datastores, the heat map will group them. You can further see the members. The larger the datastore, the larger the box.

  The **Shared Datastores** table lists the shared datastores. The table provides a summary, displaying all the datastores at a glance. They are grouped by data center. By default, the table is sorted by the least capacity remaining. There are three reclamation opportunities: powered off VM, snapshot, and orphaned VMDK.
- **Datastore Analysis**
  Select a datastore from the **Summary** table. The capacity details are automatically displayed. A snapshot that lasts beyond a few days should be investigated. Orphaned VMDK are the ones that are not associated to any VM.
  For disk space, the total capacity, allocated capacity, and the actual capacity used are displayed.
- **VM Analysis**
  To analyse at the VM level, review the **VMs in the datastore** table. Click on the VM you want to investigate further to see usage over time.
- **Local Datastores**
  The **Local Datastores Capacity** table appears at the end of the dashboard. Avoid running VMs on local data stores, unless the storage requirements can be met with a local disk and does not need vMotion.

#### Points to Note

If the underlying LUN is also thin provisioned, add visibility into the physical array. The dashboard does not have datastore clusters. If your environment uses datastore clusters, modify this dashboard or create a new one. In a large environment with many datastores and datastore clusters, add a View List to list the datastore clusters, so you get summary information. Alternatively, create a heat map, listing the datastore clusters.

### ESXi Capacity Dashboard

The **ESXi Capacity** dashboard supports the **Cluster Capacity** dashboard and is also required for the non-clustered ESXi.

#### Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

#### How to Use the Dashboard

The **Summary** heat map provides an overall view of the ESXi Host capacity, grouped by their clusters.

- Each ESXi host is represented by a box, displaying their capacity remaining.
- The ESXi host size is made constant for ease of use. If your ESXi sizes are not standardized, consider using the number of physical cores or Total CPU GHz to display the difference in sizes. Ensure that the smallest ESXi is not too small.
- Wastage is displayed by a new color. Dark gray indicates wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

The **ESXi Hosts Capacity** widget lists all the ESXi hosts in your environment, grouped by their parent cluster.

- The standalone ESXi are displayed at the bottom under No Group.
- In a large environment with many data centers, you can zoom into a specific vCenter or data center. You can also filter or search for specific ESXi hosts matching certain names.
- The **99th Percentile Performance** column takes the 99th percentile value of the ESXi Performance (%) metric. To rule out the outlier, the worst performance (which is equivalent to 100th percentile) is not considered. Also, the performance threshold is set to be stringent.

Select one of the ESXi hosts from the **ESXi Hosts Capacity** widget. All the three line charts automatically display the trend of selected ESXi host.

- Displays both total and usable utilization in terms of RAM and CPU.
- Utilization is displayed for three months and not one week. The daily average is displayed and not the hourly average and the focus is on RAM consumed and not RAM active.

#### Points to Note

- Add a drill-down to the **ESXi Capacity** dashboard. A logical place to initiate this drill-down is in the **Cluster Capacity List** widget. Link this widget into the table of ESXi host in the destination dashboard.
- A technology refresh is often used to address the capacity shortage. Consider adding a property widget that displays the hardware model and specification to help you determine the age of the hardware.

### VM Capacity Dashboard

The **VM Capacity** dashboard helps you analyze the capacity of all the VMs with the ability to analyze each VM.

#### Design Considerations

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

#### How to Use the Dashboard

- **Overall Analysis**
  The **Datacenters** table lists all your data centers. vSphere World is included so you can see all the VMs from all the data centers. Unlike infrastructure objects, there are potentially tens of thousands of VMs. The charts will take longer to refresh if you select vSphere World.

  The **VMs by Capacity** widget, groups the VMs by capacity remaining, while the **VMs by Time Remaining** widget groups the VMs by time remaining. Ideally, you want all of them to be low on capacity remaining, but high on time remaining.

The heat map provides an additional view by grouping them by cluster. The heat map helps you identify which cluster is at risk (where most of the VMs need more capacity) and which cluster provides extra resources (where most of the VMs are not using their capacity).

Review the **VMs by Capacity Remaining** heat map. The heat map provides the next level of detail by grouping the VMs by clusters, so you can see which clusters need attention. The VM size is standardized for better visualization.

- **VM Analysis**
  Review the **VMs Capacity in the Selected Datacenter** table that lists all the VMs in the selected data center. The list is sorted by the VM with the least capacity remaining. You can also sort by Time Remaining.

  Select a VM from the table. The capacity details are displayed. You can view both the CPU and memory trend over time. Three months data is displayed and is averaged to hourly so you can see the overall trend.

  Disk details are displayed using the Guest OS partition. Avoid using VM virtual disk as there may not be a direct mapping to the actual partition.

  The recommendation for right-sizing is displayed for both CPU and memory. Unlike a physical server, it is important to right-size a VM. For CPU, the CPU Usage counter is used instead of Demand. For disk, Guest OS partition level is displayed. There is no overall capacity at the VM level because different partitions have different capacity.
- **Relevant Configuration**
  Relevant configuration is automatically displayed to provide context to the VM. Information such as VM owner and business units can be useful in the analysis.

### Reclamation Dashboard

The **Reclamation** dashboard helps you manage various types of reclamation that can be carried out on VMs and datastore. It is designed for both the Capacity team and the Operations team.

#### How to Use the Dashboard

- **Overall Analysis**
  The scoreboard provides a summary of the total reclamation.

  You can select either a data center, a cluster, or a datastore. The datastore level is required as orphaned disks do not have VM association, hence it is not related to any cluster. The **Datastore** table only drives the snapshot table.

  **Powered-off VMs Distribution Size**, **Idle VMs Distribution by Memory Footprint**, and **Snapshots Distribution by Size** charts display summary information.

  Adjust the bucket size in the charts to suit your operational requirements. The reclamation potentials are presented as three bar charts, each corresponds to an area you can reclaim:
  - Powered off VMs that are no longer needed contribute to wasted disk usage. Consider deleting them to free up space or moving them to archival storage.
  - Idle VMs are running, but not being used actively. These VMs consume memory that may be used by active VMs. Consider removing these VMs to reduce memory contention.
  - Snapshots are meant to be temporary and can cause performance issues and waste disk space if not deleted after a few days.

  When reviewing each of the three tables observe that they are sorted by the largest reclamation opportunities. This allows you to get the greatest benefit from the least amount of effort. For example, focus on snapshots first, as it does not involve changing the VM. For the powered off VMs, you may want to consider VMs that have been powered off longer and these are more likely to be unneeded. Idle VMs can be challenging to reclaim since they are still running, so you might prioritize powered off VMs before attempting to reclaim idle VMs.

  The **99P CPU Usage** column displays the CPU usage at the 99th percentile during the time period. It is an easy way to check if it is indeed idle.

- **VM Analysis**
  To analyze VMs for reclamation opportunities, select a VM from one of the three tables (Powered Off VM, Idle VMs, or VM Snapshots). The selected VM will populate the widgets with the following details:
  - Powered off over time displays the amount of time a VM has been powered off.
  - CPU Usage over time provides insights into the aggregate CPU usage, including peak usage periods. This way you can validate that an idle VM has not had any brief usage.
  - VM Snapshot over time provides you with an understanding of the age and growth of snapshots on the VM. Watch for fast growing snapshots, as these can quickly consume disk space.
  - Contexts of selected VM is a summary of the VM configuration information.

**Points to Note**

To organize your reclaim efforts, it is helpful to create custom groups to make it easier to filter by department or VM owner. This can make it easier to seek approvals and communicate with anyone who may be impacted.

### vSAN Capacity Dashboard

The **vSAN Capacity** dashboard complements the vSphere **Cluster Capacity** dashboard by displaying capacity related to vSAN. To manage vSAN capacity, use both dashboards.

**Design Considerations**

The dashboard focuses on vSAN specific metrics, but does not list non-vSAN clusters.

See Capacity Dashboards for common design considerations among all the dashboards for capacity management.

**How to Use the Dashboard**

- **Overall Analysis**
  The **Clusters by Capacity Remaining** and **Clusters by Time Remaining** bar charts focus on vSAN disk space and not compute and network.

  Use the **vSAN Clusters** table to view vSAN specific metrics.
- **Cluster Analysis**
  Select a vSAN cluster from the **vSAN Clusters** table. The detailed capacity is automatically displayed.

  The **Utilization** widget displays the utilization for all three elements, as you need to consider all three. Network is not shown as typically it is not a problem.

  Like a physical array, there can be hot spots and imbalance. The **Is Disk Group Space Utilization Balanced** heat map displays individual disk groups.

  Reclaimable storage is a key component of proactive capacity management. You can view details for both VMs and non-VMs.

  You can view the **Dedupe and Compressed** scoreboard for more details in this area.
- **Disk Group Analysis**
  If there is imbalance, you can analyze each disk group. The **Disk Groups Selected in Selected vSAN Cluster** displays all the disk groups in the cluster. Their usage may not be similar, but should not deviate drastically. To view the disk group usage trend, click on a disk group.
- **VM Analysis**
  You can analyze individual VMs in the selected cluster from the **VMs in the selected vSAN Cluster** table and check their usage and snapshot. To view the trend in Usage, click on a VM. In addition, the relevant configurations of the VM is displayed.

### vSAN Stretched Clusters

The vSAN Stretched Clusters dashboard provides an overview of the cluster resources used across vSAN fault domains. Using the stretched clusters dashboard you can monitor the resource consumption at the site level for Preferred Sites and Secondary Sites. You can create custom dashboards for specific vSAN stretched cluster metrics.

**Where to View vSAN Stretched Cluster Objects**

From the left menu, click **Operations › Dashboards**. From the Dashboards panel, click **All › Capacity › vSAN Stretched Clusters**.

You can also view the vSAN stretched cluster objects from **Inventory › VMware vSAN › vSAN and Storage Devices › vSAN Clusters**, if the vSAN cluster is a stretched cluster.

The vSAN Stretched Clusters dashboard provides information about CPU Capacity, Cores, Memory Capacity, and Disk Capacity for the Preferred Site and the Secondary Site. You can identify the vSAN stretched clusters running out of capacity looking at the utilization metrics.

### vSAN ESA Capacity Dashboard

The vSAN ESA ( Express Architecture) Capacity dashboard provides an overview of the capacity available across all the vSAN clusters. The dashboard displays the time remaining before capacity is projected to run out, the amount of capacity remaining, the number of vSAN clusters that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

**How to Use the Dashboard**

- **Overall Analysis**
  The **All Clusters by Capacity Remaining** and **All Clusters by Time Remaining** bar charts focuses on the number of vSAN disk against the used disk capacity percentage across vSAN clusters.
- **Cluster Analysis**
  Select a vSAN ESA cluster from the **vSAN Clusters** table. The detailed capacity is automatically displayed.

  The **Utilization** widget displays the utilization for all three elements, as you need to consider all three. Network is not shown as typically it is not a problem.

  The **Is Storage Pool Space Utilization Balanced** heat map displays individual storage pools.

  Reclaimable storage is a key component of proactive capacity management. You can view details for both VMs and non-VMs.
- The **Storage Pools in Selected vSAN Cluster** widget displays the utilization for Disks and Disk Space.
- Like a physical array, there can be hot spots and imbalance. The **Is Physical Disk Space Utilization Balanced** heat map displays individual physical disks.
- **VM Analysis**
  You can analyze individual VMs in the selected cluster from the **VMs in the selected vSAN Cluster** table and check their usage and snapshot. To view the trend in Usage, click on a VM. In addition, the relevant configurations of the VM is displayed.

### vSAN OSA Capacity Dashboard

The vSAN OSA (Original Storage Architecture) Capacity dashboard provides an overview of the capacity available across all the vSAN clusters. The dashboard displays the time remaining before capacity is projected to run out, the amount of capacity remaining, the number of vSAN clusters that might fit in the remaining capacity, and reclaimable resources that can increase the available capacity.

**How to Use the Dashboard**

- **Overall Analysis**
  The **All Clusters by Capacity Remaining** and **All Clusters by Time Remaining** bar charts focuses on the number of vSAN disk against the used disk capacity percentage across vSAN clusters.
- **Cluster Analysis**
  Select a vSAN OSA cluster from the **vSAN Clusters** table. The detailed capacity is automatically displayed.

  The **Utilization** widget displays the utilization for all three elements, as you need to consider all three. Network is not shown as typically it is not a problem.

  Like a physical array, there can be hot spots and imbalance. The **Is Disk Group Space Utilization Balanced** heat map displays individual disk groups.

  Reclaimable storage is a key component of proactive capacity management. You can view details for both VMs and non-VMs.
- **Disk Group Analysis**
  If there is imbalance, you can analyze each disk group. The **Disk Groups Selected in Selected vSAN Cluster** displays all the disk groups in the cluster. Their usage may not be similar, but should not deviate drastically. To view the disk group usage trend, click on a disk group.
- **VM Analysis**
  You can analyze individual VMs in the selected cluster from the **VMs in the selected vSAN Cluster** table and check their usage and snapshot. To view the trend in Usage, click on a VM. In addition, the relevant configurations of the VM is displayed.

## Configuration Dashboards

As an operations management software, VMware Aria OperationsVMware Cloud Foundation Operations focuses on the impact to day-to-day operations a product has, rather than the feature of the product itself. Products under monitoring, such as vSphere and vSAN, can have features that are related, but have a different impact on operations. For example, vSphere provides Limits, Reservation, and Shares for the VM.

Limits, Shares, and Reservation. As a feature, they are closely related, and appear in the same dialog box and must be learned as one. However, they impact operations differently. The following table describes that in more detail.

| | | | |
|---|---|---|---|
| VM Limit | Impacts the VM | • Should not be used. Right-size instead.<br>• Results in unpredictable performance of the Guest OS. | Incorrect Configuration |
| VM Reservation | Impacts the Infrastructure | • Keep the total amount low, and relative to the total capacity of the cluster.<br>• Absolute value. A 2-GHz reservation is in fact a 2-GHz reservation.<br>• Results in suboptimal infrastructure capacity, as overcommit is not possible. | Suboptimal Configuration |
| VM Share | Impacts the Infrastructure | • Keep the number of variations to below three. One for each class of service.<br>• Relative value. 2000 worth of reservations depends on the value of other VM reservations. Be careful when you move the VM to another cluster, as the relative value changes.<br>• Results in complex operations. It is harder to troubleshoot performance when the dynamic entitlements of each VM fluctates more. | Complex Configuration |

VMware Aria OperationsVMware Cloud Foundation Operations follows the principle that there are different impacts on operations and applies a methodology for looking at configuration. It does not group the settings by features or objects. Rather, it begins with the impact and prioritizes what can be done.

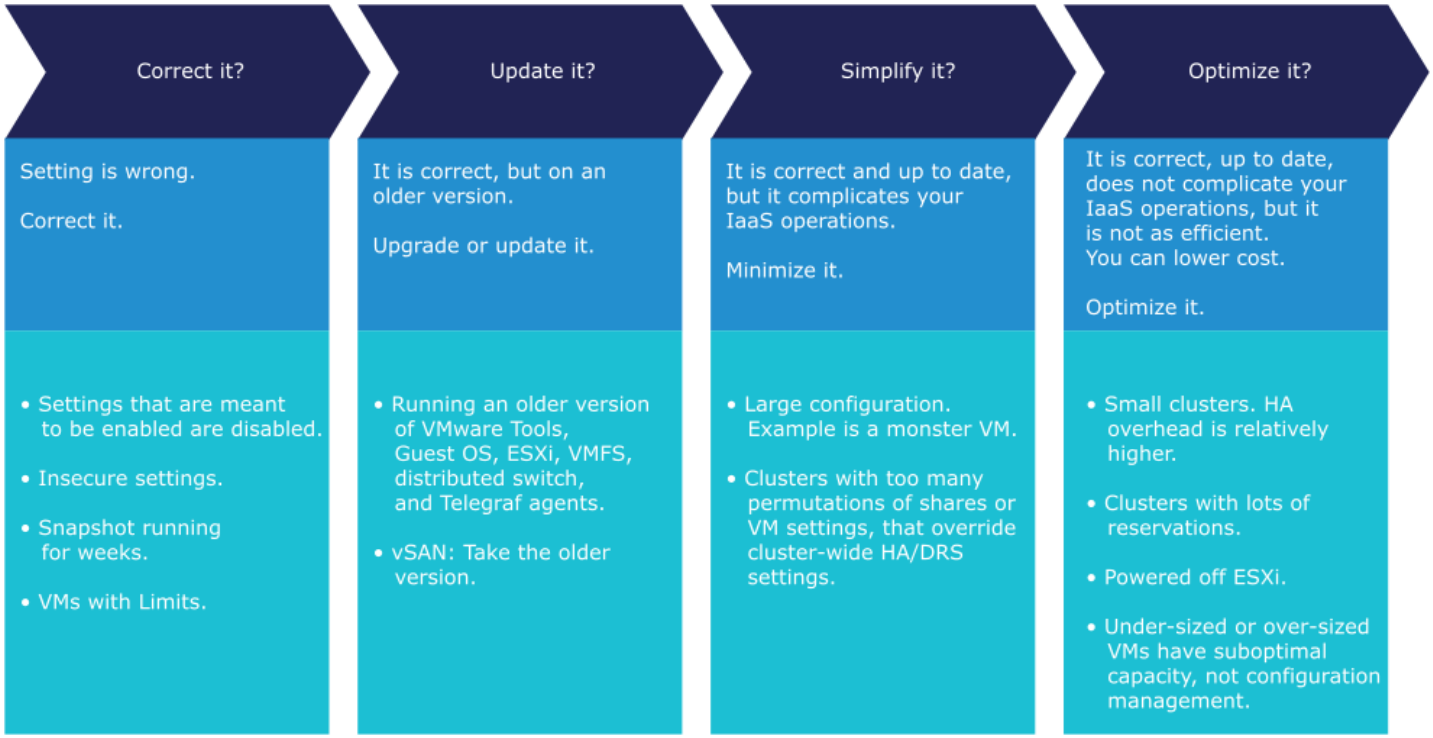


Each operation is unique and as a result, customers run operations differently. What is right for other customers, might not be right for you. Even in the same environment, what is right for a development environment might not be appropriate for a production environment.

The following table lists some of the areas for improvement for the operations in your environment:

## Areas of Improvement

| | Correct it? | Update it? | Simplify it? | Optimize it? |
|---|---|---|---|---|
| **IaaS Consumer:**<br>• Process<br>• Applications<br>• Guest OS<br>• Container<br>• VM | • Java JVM or Database \ memory config too large relative to Guest OS<br>• Guest \ Metric not collecting<br>• Guest \ High TX Broadcast packets<br>• VM \ Tools not installed<br>• VM \ Tools not running<br>• VM \ CPU Limit<br>• VM \ Memory Limit<br>• VM \ Old Snapshot<br>• VM \ On local Datastore | • Guest OS \ Tools<br>• Guest OS \ Windows<br>• Guest OS \ Linux<br>• Guest OS \ Telegraf agent<br>• VM \ Hardware (vmx) | • VM \ Large VM (CPU, RAM, Disk)<br>• VM \ lots of disks, NIC card<br>• VM \ lots of IP address.<br>• VM \ with RDM<br>• VM \ on multiple datastores<br>• VM \ Fault Tolerant<br>• VM \ SRM protected<br>• VM \ Hot Add\|Remove \ CPU<br>• VM \ Hot Add\|Remove \ RAM | • Java JVM or Database \ memory config too small relative to Guest OS<br>• Guest OS \ no visibility<br>• Container \ smaller than the parent VM<br>• VM \ Tools unmanaged<br>• VM \ bigger than the whole ESXi cores.<br>• VM \ bigger than CPU socket.<br>• VM \ Large Snapshot<br>• VM \ Reservation. |
| **IaaS Provider:**<br>• Telegraf<br>• ESXi<br>• Cluster<br>• Datastore & Cluster<br>• Switch and Port Group<br>• Hardware<br>• NSX<br>• vSAN | • ESXi \ vMotion disabled<br>• ESXi \ Disconnected from vCenter<br>• ESXi \ Maintenance Mode<br>• ESXi \ NTP disabled<br>• ESXi \ Standalone<br>• Cluster \ Admission Control disabled<br>• Cluster \ HA disabled<br>• Cluster \ HA Failover %<br>• Cluster \ DRS disabled<br>• Cluster \ DRS manual<br>• Cluster Inconsistency<br>  • BIOS, ESXi: version<br>  • BIOS, ESXi: Power Management<br>  • ESXi Storage Path<br>  • ESXi Hardware<br>• Datastore Cluster inconsistency<br>  • Capacity<br>  • Performance<br>• Datastore \ single path<br>• Datastore \ no path. This is unlikely.<br>• NSX \ no redundancy for Controller, Manager | • ARC \ server<br>• ARC \ agent<br>• ESXi \ hardware<br>• ESXi \ vSphere<br>• ESXi \ 1 Gb NIC.<br>• Server \ not on warranty<br>• vCenter \ version<br>• Datastore \ VMFS version<br>• vSAN \ version<br>• Switch \ version<br>• NSX \ version | • ESXi \ Too many variations. No standard<br>• Cluster \ Many VM Shares (CPU)<br>• Cluster \ Many VM Shares (RAM)<br>• Cluster \ Resource Pools<br>• Cluster \ Stretched compute + storage<br>• Cluster \ 32 nodes or more<br>• Cluster \ VM to Host affinity<br>• Cluster \ Too many storage paths<br>• Datastore \ Shared by >1 cluster WLP uses this<br>• Datastore \ Many paths<br>• Network \ LBT?<br>• Network \ MAC Address change | • ESXi \ low CPU cores count<br>• ESXi \ low RAM size<br>• ESXi \ Powered Off<br>• ESXi \ HT Disabled<br>• ESXi \ 4 socket or higher.<br>• Cluster \ small clusters \ host especially for vSAN<br>• Cluster \ small clusters \ CPU<br>• Cluster \ small clusters \ RAM<br>• Cluster \ EVC Mode<br>• Cluster \ High Reservation<br>• Cluster \ DRS Automation Level<br>• Cluster \ DPM disabled<br>• vSAN \ All Flash: Dedupe disabled<br>• vSAN \ All Flash: Compressed disabled<br>• Datastore \ small<br>• Datastore \ low VM count<br>• Datastore \ no ESXi<br>• Distributed Switch \ unused |

## Design Considerations

The dashboards display configurations that need immediate attention, before displaying the overall configuration. This helps you take measures toward optimizing configuration.

Operations vary among customers, and as a result, it is not possible to design one dashboard to meet every customer's operational needs. A configuration that is important for one customer might not be relevant for another customer. Tailor the dashboard to your unique environment. You can collapse or expand the widgets to allow relevant data to be displayed.

The overall layout is designed to balance ease of use, performance (loading time of the dashboard page), and completeness of configuration check. As a result, not all configuration settings are displayed. Lack of screen real estate is another consideration behind the design.

### Cluster Configuration Dashboard

Use the **Cluster Configuration** dashboard to view the overall configuration of vSphere clusters in your environment, especially configurations that need attention.

## Design Considerations

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

## How to Use the Dashboard

The **Cluster Configuration** dashboard is organized into sections for ease of use.

- The first section of the dashboard consists of three bar-charts. They correspond to the three main features of vSphere clusters, namely High Availability (HA), Dynamic Resource Scheduler (DRS), and Distributed Power Management (DPM).

  - HA: The best practice is to activate HA admission control. You can specify the admission control policy in the vCenter and the threshold for failover shares.
  - DRS: The best practice is to activate DRS. Envision the vSphere cluster as a single logical computer that balances within itself.
  - DPM: The best practice is to activate DPM in an environment where environmental concern is the top priority or the high peak rarely occurs as most of the time you run very low utilization.
- The second section of the dashboard consists of eight pie-charts. They show the relative distribution of key configurations.
  - Two of the bar-charts cover admission control. You must activate admission control. The pie-charts displays the policy code instead of the policy name, as it is based on the property: `Cluster Configuration | Das Configuration | Active Admission Control Policy`. The mapping between the code to name is:
    - -1 = Disabled
    - 0 = Cluster Resource Percentage
    - 1 = Slot Policy (Powered-on VMs)
    - 2 = Dedicated Failover Hosts
  - There are two bar-charts that cover the HA Failover Share. One for CPU, and one for memory.
  - The next two bar-charts cover DRS settings. You might want to fully automate DRS, which means that there is no operator intervention required for both initial VM placement and subsequent load balancing, but with a moderate migration threshold (value = 3.0). The value ranges from 1.0 to 5.0.
  - There are two pie-charts that show reservation. One for CPU and one for memory. Minimize the total reservation value as it prevents overcommit of resources and hence results in less optimal utilization. Memory reservation can remain and occupy the memory space of the ESXi host, even though the VM does not use the memory anymore. Consider the analogy of unused files that you have not opened for months in the c:\ drive of your laptop. They still take up space on the hard disk. Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service.
- The third section of the dashboard consists of two bar-charts. They show the absolute distribution of the clusters.
  - The first bar-chart displays the cluster grouped by the number of ESXi hosts. Small clusters, defined as having a lower number of ESXi hosts, have a higher overhead while large clusters have a higher risk if there are cluster-wide outages. Performance risk is lower, because there are more nodes that DRS can tap on, but if there is an actual problem, troubleshooting can be tougher, because there are more nodes to analyze. For large clusters, have a disaster recovery plan as an unexpected cluster-wide outage can impact many VMs.
- The fourth section of the dashboard lets you drill-down to an individual cluster.
  - A table lists all the clusters with their key configuration. You can export this list as a spreadsheet for further analysis or reporting.
  - Select a cluster. The list of ESXi hosts under the cluster, with shares and resource pool information, is automatically filled up.
  - Keep the number of distinct shares to below three (or at a minimum), matching the distinct classes of service. Avoid providing different services to individual VMs as that increases the complexity of the cluster performance.
  - Keep the number of resource pools minimal.
  - Some of the columns are color coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

**Points to Note**

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- `No data to display` does not imply that there is something wrong with data collection by VMware Aria OperationsVMware Cloud Foundation Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

### ESXi Configuration Dashboard

Use the **ESXi Configuration** dashboard to view the overall configuration of the ESXi hosts in your environment, especially the configurations that need attention.

**Design Considerations**

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

**How to Use the Dashboard**

The **ESXi Configuration** dashboard is organized into sections for ease of use.

- The upper section of the dashboard displays basic ESXi configurations that should be standardized for ease of operations.
  - There are six pie-charts that are displayed as one set because there is a relationship between their values. There should be a correlation between them. Ideally, the ESXi version, the ESXi build, and the BIOS must be identical across all ESXi hosts in a cluster. Keep the variations of the hardware model, NIC speed, and storage path minimal. The more complex the pie-chart, the more variants you have. This results in complex operations, that potentially results in higher operating expenses.
  - The configurations should reflect your current architecture standard. Each pie-chart counts the occurrences of a particular value. A large slice signifies that the value is the most common value, and if that is not your current standard, then you must address it.
- The second section of the dashboard displays configurations that are potentially suboptimal.
  - The three bar-charts display various size dimensions of the ESXi hosts. The bar-charts are designed to be seen as one set. Ensure that there are a minimal number of variations to reduce complexity.
  - Smaller ESXi hosts have a relatively higher overhead, and are limited in running larger VMs. If they have a low core count, they might be using an outdated CPU. Small ESXi hosts are more expensive on a per core, per GB, per rack unit basis than larger ones if they occupy the same space. However, a 4-CPU socket ESXi host is likely to be too large, resulting in a concentration risk (too many VMs in a single ESXi host). Maintain a good balance that balances your budget and risk constraints.
  - Adjust the distribution chart bucket size to fit your environment.
- The third section of the dashboard displays configurations that you might want to avoid.
  - The six bar-charts focus on security, availability, and capacity settings that you can set as a standard. For example, you should consider activating the NTP daemon for a consistent time, which is critical for logging and troubleshooting.
  - The three tables list the actual ESXi hosts that are in a non-productive state. They can be on maintenance mode, powered off, or in a disconnected state.
- The last section of the dashboard displays all the ESXi hosts in your environment.
  - You can sort the columns and export the results into a spreadsheet for further analysis.

  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.

**Points to Note**

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- `No data to display` does not imply that there is something wrong with data collection by VMware Aria OperationsVMware Cloud Foundation Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result, there is nothing to display.
- In a large environment, create a filter for this dashboard. Group by the class of services such as, gold, silver, and bronze. Default the selection to gold. In this way, your monitoring is not cluttered with less critical workloads.
- For complete visibility, consider adding physical server monitoring by using the appropriate management pack. For more information, see the following page.

### Network Configuration Dashboard

Use the **Network Configuration** dashboard to view the overall configuration of vSphere distributed switches in your environment, especially for the areas that need your attention.

**Design Considerations**

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

**How to Use the Dashboard**

The dashboard is organized into two sections for ease of use.

- The first section displays network configurations that need your attention.
  - There are five bar-charts that focus on critical security settings.
  - The last bar-chart displays the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.
- The second section provides overall configuration information, with the ability to drill down to a specific switch.
  - Click the row to select a switch from the list.
  - The ESXi hosts, port groups, and the VMs on the switch are displayed.
  - Review each of the tables. For the ESXi host table, ensure that the settings are consistent.
  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.
  - You can sort the columns and export the result into a spreadsheet for further analysis.

**Points to Note**

- `No data to display` does not imply that there is something wrong with data collection by VMware Aria OperationsVMware Cloud Foundation Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- For complete visibility, consider adding physical network device monitoring by using the appropriate management pack. For more information, see the following page.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

### VM Configuration Dashboard

Use the **VM Configuration** dashboard to view the overall configuration of virtual machines in your environment, especially for the areas that need attention.

**Design Considerations**

See the Configuration Dashboards page for common design considerations among all the dashboards for configuration management.

As there are many configurations to be verified, if you have a larger screen, add additional checks as you deem fit, or add legends to the pie-charts.

**How to Use the Dashboard**

- Click the row to select a data center from the data center table.
  - In a large environment, loading thousands of VMs increases the web page loading time. As a result, the VM is grouped by data center. In addition, it might make sense to review the VM configuration per data center.
  - For a small environment, vSphere World is provided, so you can view all the VMs in the environment.

  The **VM Configuration** dashboard is organized into three sections for ease of use. All the three sections display the VM configuration for the selected data center.
- The first section covers limits, shares, and reservations.
  - Their values can easily become inconsistent among VMs, especially in an environment with multiple vCenter Servers.
  - Shares should be mapped to a service level, to provide a larger proportion of shared resources to those VMs who pay more. This means that you should only have as many shares as your service levels. If your IaaS provides gold, silver, and bronze, then you should have only three types of shares.
  - Value of the shares and reservation is relative. If you move a VM from one cluster to another (in the same or different vCenter), you might have to adjust the shares.
  - Reservation impacts your capacity. Memory reservation works differently from CPU reservation, and it is more permanent.
- The second section covers VMware Tools.
  - VMware Tools is a key component of any VM, and should be kept running and up to date.
- The third section covers other key VM configurations.
  - Keep the configurations consistent by minimizing the variants. This helps to reduce complexity.
  - **VM Network Cards** widget. If you suspect that your environment might have a VM with no NIC, consider adding it as a dedicated bucket.
- The last section of the dashboard is collapsed by default.
  - You can view all the VMs with their key configurations.
  - You can sort the columns and export the results into a spreadsheet for further analysis.

**Points to Note**

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- `No data to display` does not imply that there is something wrong with data collection by VMware Aria OperationsVMware Cloud Foundation Operations. It might signify that none of the objects meet the filtering criteria of the widget, and as a result there is nothing to display.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.
- The pie-chart and bar-chart cannot drive other widgets. For example, you cannot select one of the pie-slices or buckets, and expect it to act as a filter to a list or a table.
- You can apply a specific color in a pie-chart or distribution chart for a specific numeric value, but not string value. For example, you cannot apply the color red to the value `Not Installed`.

**vSAN Configuration Dashboard**

The **vSAN Configuration** dashboard provides overall configuration details and is useful in large clusters with many vSANs, where you have to follow a certain standard configuration.

**Design Considerations**

See Configuration Dashboards for common design considerations among all the dashboards for configuration management.

**How to Use the Dashboard**

The **vSAN Configuration** dashboard is organized into three sections for ease of use.

- The first section displays six pie-charts.
  - There are five bar-charts that focus on critical security settings.
  - The last bar-chart shows the version of the vSphere Distribution Switch. Aim to keep the version current, or match your vSphere version.
- The second section displays three bar-charts.
  - The three bar-charts together provide a good overview of the vSAN key capacity configuration. By analyzing the distribution, you can identify if you have capacity configuration that is outside your expectation.
- The last section of the dashboard displays all the vSAN clusters with their key configuration.
  - Some of the columns are color-coded to facilitate quick reviews. Adjust their threshold to either reflect your current situation or your desired ideal state.
  - You can sort the columns and export the result into a spreadsheet for further analysis.

**Points to Note**

- The number of buckets in the pie-chart or bar-chart are balanced between the available screen estate, ease of use, and functionality. Modify the buckets to either reflect your current situation or your desired ideal state.
- To view the content of a slice in a pie-chart or a bucket in a bar-chart, click on it. The list cannot be exported. Clicking an object name, takes you to the object summary page. The page provides key configuration information, with other summary information.

### Workload Management Configuration Dashboard

This dashboard provides a quick configuration summary of all the key objects associated with workload management such as Supervisor Clusters, Namespaces, vSphere Pods and Tanzu Kubernetes clusters. It is essential that the configuration is consistent across all the objects. Configuration drifts may result in inconsistent performance or availability of the applications leveraging workload management Kubernetes constructs.

Use the dashboard to ensure that the configuration is consistent across all objects.

You can view the following widgets in the dashboard.

- **Environment Summary**
- **Supervisor Cluster Versions**
- **Cluster Status**
- **Pod Data**
- **Supervisor Cluster Configuration Summary**
- **Pod Configuration Summary**
- **Kubernetes cluster Configuration Summary**
- **Namespace Configuration Summary**

### Consumer \ Correct it? Dashboard

The **Consumer \ Correct it?** dashboard complements the main VM configuration dashboards by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

## Design Considerations

The dashboard is designed to focus on VMs that need attention. Lists are used to keep it simple, and show actual objects. The lists can be tailored using the filter and the custom group. The lists can also be exported for an offline discussion.

The dashboard is extendable, reflecting the reality that different customers have a different set of settings to verify. Since the dashboard layout is a collection of tables (List View), you can extend it by adding more tables. You can add more List View widgets to verify the VM configurations that your operations require.

## How to Use the Dashboard

The **Consumer\Correct it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Tools Widgets:
  - Using VMware Tools has multiple benefits. For the list of benefits, refer to KB 340.
  - VMware Aria OperationsVMware Cloud Foundation Operations uses VMware Tools to retrieve Guest OS metrics. Without this, right-sizing VM memory can be inaccurate, because the hypervisor metrics (VM Memory Consumed and VM Memory Active) are not designed to measure Windows or Linux memory utilization. ESXi VMkernel does not have visibility into the Guest OS for security reasons.
  - Independent software vendor (ISV) support is the most common reason that VMware Tools is not installed. The ISV vendor might claim that no additional software is installed in their appliance unless they have certified it. For more information about VMware Tools, see the VMware Tools documentation.
  - If VMware Tools is installed, there might be reasons why the application team deactivates it. The Infrastructure team should inform and educate their application team, and document the technical recommendations about why VMware Tools is recommended to be running all the time.
- CPU Limits and Memory Widgets:
  - It is recommended that you do not use memory and CPU limits as it can result in an unpredictable performance. The Guest OS is not aware of this restriction as it is at the hypervisor level. It is recommended that you shrink the VM instead.
- Guest OS Counters Missing Widget:
  - There is no visibility into the Guest OS performance counters because the requirements are not met. The memory counter is especially important as VM Consumed and VM Active are not replacements for Guest OS counters. See KB 55675 for more details.
- Old Snapshot Widget:
  - Ensure that the snapshot is removed within one day after the change request. If not, it might result in a large snapshot and impact the performance of the VM.

## Points to Note

- Add a banner summary to the top of this dashboard so that you can verify if there is an incorrect confirmation. Add a scoreboard and select the World object and then collapse all the tables below. Create a super metric for each summary and apply it to the World object.
- In a large environment, create a filter for this dashboard to help you to focus on a segment of the environment. Group it by a class of service such as, gold, silver, and bronze. Default the selection to gold, your most important environment. In this way, your monitoring is not cluttered with less critical workloads.
- There are other VM configurations that maybe relevant to your environment. Review the list of VM settings that you might want to add to this dashboard.
- For context, add a property widget that lists the selected VM properties. In this way, you can check the property of your interest without leaving the screen. Multiple List View widgets can drive the same property widget, so you do not have to create one property widget for each List View.
- If your operations require it, add a list of VMs that do not have these three key performance counters: CPU Run Queue, CPU Context Switch, and Disk Queue Length.

### Consumer \ Optimize it? Dashboard

The **Consumer \ Optimize it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Optimize it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities. A suboptimal configuration might not impact performance or increase complexity, but it can be more expensive.

#### Design Considerations

The **Consumer \ Optimize it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

#### How To Use the Dashboard

The **Consumer \ Optimize it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- VM Reservation:
  - VM reservation causes a positive impact on the VM, but a negative impact on the cluster. Total reservation cannot exceed cluster capacity. This creates a suboptimal cluster as VMs do not use the entire assigned memory at the same time.
  - VM reservation places a constraint on the DRS placement and HA calculation. Avoid using reservation as a means to differentiate performance SLA among all the VMs in the same cluster. It is difficult to correlate CPU Ready with CPU Reservation. A VM CPU Ready does not improve two times because you increase its CPU reservation by two times. There is no direct correlation.
- Guest OS visibility:
  - Since your workloads are sharing resources and are over-committed, your operations are easier if you know what is running inside. This helps with monitoring and troubleshooting, resulting in optimal operations.
  - For critical VMs, consider logging the Guest OS, such as Windows and Linux, to capture errors that do not surface as metrics. These errors typically appear as events in the log files or in the event database in the case of Windows. Use VMware Cloud Foundation Operations for logs to parse Windows events into log entries that can be analyzed.
- Snapshot:
  - Old snapshots tend to be larger. They consume more space and have a higher chance of impacting performance.

#### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? dashboard. This dashboard follows the same design considerations, and as a result shares limitations and customization ideas.

### Consumer \ Simplify it?

The **Consumer \ Simplify it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Simplify it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

#### Design Considerations

The **Consumer \ Simplify it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

### How to Use the Dashboard

The **Consumer \ Simplify it?** dashboard is a collection of tables (List View), that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Large VMs (CPU, Memory, and Disk):
  - A large VM, relative to the underlying ESXi host and datastore, requires more careful planning (Day 0) and monitoring (Day 2).
  - Ensure that the VM size does not exceed the size of the underlying ESXi host. If your ESXi host has CPU hyper-threading, do not count the logical processor. Instead, count the physical core. For best performance, keep it within a (non-uniform memory access) NUMA boundary.
  - During monitoring, verify if the VM is highly utilized. If the VM vCPU count is equal to the ESXi cores, and the VM is running at almost full capacity, you might not be able to run other VMs. Large VMs can impact the performance of other VMs, especially if it is given higher shares. Only when the large VM is under-utilized, can the ESXi hosts run other VMs.
  - If the number of configured vCPUs on a VM is higher than the number of cores per socket on the ESXi, the VM can experience the NUMA effect. If the ESXi has more than one physical CPU (socket), cross-NUMA access negatively impacts performance.
  - The larger the VM, the longer the time required to vMotion, Storage vMotion, and backup.
  - For disk space, if the disk is thin-provisioned and under-utilized, you can deploy other VMs in the same datastore. Ensure that the snapshot is tracked closely, as the risk of capacity running out is higher for a large virtual disk.
- VMs with many virtual disks:
  - It is simpler to have a 1:1 mapping between Guest OS partitions and the underlying virtual disk (VMDK or RDM).
  - For performance and capacity, evaluate the disks and partitions. Each virtual disk must be monitored in terms of IOPS, throughput, and latency. Having multiple virtual disks increases the monitoring and troubleshooting need.
  - If the reason for having many virtual disks is performance, identify which counter serves as proof that multiple virtual disks are required. It is possible that the performance required is met by a single virtual disk.
- VM with many IP addresses or NICs:
  - A VM might need multiple networks, such as production, back up, and management. It is recommended that you route the network interfaces through the NSX-Edge VM. A VM that has multiple network interfaces can bridge the network, causing security risks or network problems.
  - A VM that is part of multiple networks can do so with just a single NIC. A single NIC can be configured to access multiple networks, with each interface having their own IP configuration.

### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

## Consumer \ Update it? Dashboard

The **Consumer \ Update it?** dashboard complements the main VM configuration dashboard by displaying the actual VMs, with their relevant information. The dashboard is designed for vSphere administrators and the platform team, to facilitate follow-up action with the VM owners. The **Consumer \ Update it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

### Design Considerations

The **Consumer \ Update it?** dashboard follows the same design considerations specified for the Consumer \ Correct it? dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

### How to Use the Dashboard

The **Consumer \ Update it?** dashboard is a collection of tables (List View), which can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated Tools Widget:
  - Lists all the VMware Tools versions that are still supported. Tailor the filter to fit your operational needs.
- Outdated VM Hardware Widget:
  - Lists all the VM vmx versions that are not 13, 14, 15, or 16. Tailor the filter to fit your operational needs.
- Outdated Windows and Red Hat Widgets:
  - Lists all the Windows client versions that are not version 10.
  - Lists all the Windows server versions that are not versions 2016 and 2019.
  - Lists all the RHEL versions that are not version 7 or 8.
  - If you run other operating systems like Ubuntu, clone the widget. You can also repurpose the widget if you do not run RHEL and Windows.

### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? dashboard. This dashboard follows the same design considerations, and hence shares limitations and customization ideas.

## Provider \ Correct it? Dashboard

The **Provider \ Correct it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Correct it?** dashboard is one of the eight dashboards that check the environment for optimization opportunities.

### Design Considerations

The **Provider \ Correct it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

### How to Use the Dashboard

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations.
  - A cluster is the smallest logical building block for compute. Consider it as a single computer with physically independent components. As a result, consistency matters.
  - Clusters with DRS set to manual. This means that DRS initiated vMotion does not take place unless the administrator manually approves it. Since DRS calculates every five minutes, your quick approval is required to prevent a change of condition.
  - Clusters with HA disabled. Without high availability provided by the infrastructure, each application must protect itself from an infrastructure failure.
  - Clusters with DRS disabled. DRS focuses on performance and capacity, while HA focuses on availability. Without DRS, you must build a buffer on every ESXi host to cope with peak demand.
  - Clusters with Admission Control disabled. Reservation is respected only when Admission Control is activated.
- The second section covers the ESXi host configurations.
  - ESXi with Network Time Protocol disabled. Logs are a critical component of operations, and are the main source of information in troubleshooting. While troubleshooting performance across objects, the sequence of logs determines which event is the likely root cause, as the oldest event starts the chain of events.
  - A disconnected ESXi host indicates that the ESXi host is not participating in HA and you cannot migrate any VM on it.

  - An ESXi host that is in maintenance mode does not contribute resources to the cluster or the data center if there is a standalone ESXi.
- The third section covers ESXi host configurations that must be consistent within a cluster.
  - BIOS version and ESXi versions.
  - BIOS Power Management, ESXi: Power Management. Ideally, should be set to OS controlled. The ESXi level should be set to balance level.
  - ESXi Storage Path. Ensure that the number of paths and the path policies are identical.
  - ESXi hardware specifications. Different specifications can result in inconsistent performances experienced by the VM.

**Points to Note**

- See the **Points to Note** section as specified in the Consumer \ Correct it? dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.
- If you have a standalone ESXi, and you plan to replace it with a clustered ESXi host, add a table to list them.
- Based on your security settings, add a table to check the Distributed Switch and Port Group to ensure that security settings such as promiscuous mode, are used correctly.

### Provider \ Optimize it? Dashboard

The **Provider \ Optimize it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Optimize it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

**Design Considerations**

The **Provider \ Optimize it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

**How to Use the Dashboard**

The dashboard is organized into three sections for ease of use.

- The first section covers vSphere cluster configurations:
  - A small cluster has a higher HA overhead when compared to a large one. For example, a three-node cluster has 33% overhead while a 10-node cluster has 10%. For vSAN, a low number of hosts limits the availability option. Your choice of FTT is relatively more limited.
  - Many small clusters result in silos of resources. As a cluster behaves like a single computer, ensure that it has enough CPU cores, CPU GHz, and Memory. For ESXi in 2020, it is typical to have 512 GB of RAM. This results in 12 TB of RAM for a 12-node cluster, which is enough for DRS to place many VMs as it balances them.
  - If there is a lot of reservation, add a list for clusters with a relatively high reservation. If your clusters are of different sizes, use a super metric to convert the reservation value to a percentage.
- The second section covers ESXi host configurations.
  - Small ESXi. A small host faces scalability limits in running a larger VM. While a 2-socket, 32-cores, 128 GB memory ESXi can run 30 vCPU, 100 GB RAM VMs, the VM experiences a non-uniform memory access (NUMA) effect.
  - ESXi powered off. You can mark the ESXi hosts for decommissioning using the custom property feature of VMware Aria OperationsVMware Cloud Foundation Operations. You can then create a separate list, so they are not overlooked.
- The third section cover storage and network.
- 
- Unused network (distributed port group). This is a potential security risk as you might not monitor it.

**Points to Note**

- See the **Points to Note** section as specified in the Consumer \ Correct it? dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.
- For CPU cores, a change in vSphere licensing means that the ideal core is 32-cores per CPU socket. This maximizes the software license. For more information, see the vSphere Pricing Model.

### Provider \ Simplify it? Dashboard

The **Provider \ Simplify it?** dashboard complements vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Simplify it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

**Design Considerations**

The **Provider \ Simplify it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

**How to Use the Dashboard**

- Click the row in the **Clusters** widget to select one of the clusters from the table.
  - A cluster is more complex to operate when it has resource pools, shares, and limits.
- Review the list of resource pools:
  - Ensure that the number of VMs in each resource pool reflects the intended settings for the VM. The resource pool value is divided and shared among the VMs. The more the VMs, the lesser the resources allotted to each VM.
  - Verify if there are VMs who are siblings to the resource pools.
  - Verify if the resource pools are further split into subresource pools.
- Review the CPU Share and Memory Shares pie-charts:
  - Multiple combinations of shares, especially both CPU and memory, makes troubleshooting difficult.
  - Each share must map to exactly one class of service, such as one for gold and one for silver as the shares define the class of service. Shares are also relative, meaning the value depends on the value of sibling objects, such as, resource pool or VM. Ensure that the values are consistent across clusters to avoid unintended consequences while moving the VM to another cluster.
- Review the CPU Reservation and Memory Reservation tables:
  - High total reservation, especially both CPU and memory, complicates the cluster operations as it impacts the HA slot calculation, and limits the DRS choice of placement.
- Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

**Points to Note**

See the **Points to Note** section as specified in the Consumer \ Correct it? dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

### Provider \ Update it? Dashboard

The **Provider \ Update it?** dashboard complements the main vSphere configuration dashboards by displaying the actual vSphere objects, with their relevant information. The dashboard is designed for vSphere administrators and the platform team. The **Provider \ Update it?** dashboard is one of the eight dashboards that checks the environment for optimization opportunities.

As part of operations best practices, keep the infrastructure up to date. Running outdated components that are too far behind the latest version, can cause support problems or upgrade problems. It is common that the fix for the problem is

only available in the later versions. Outdated hardware can also result in higher operating costs. Outdated hardware might cost more data center footprint, such as rack space, cooling, and UPS. Refreshing your technology and consolidation are two common techniques to optimize cost.

#### Design Considerations

The **Provider \ Update it?** dashboard follows the same design considerations specified in the Consumer \ Correct it? dashboard. The eight Configuration > Review dashboards form an optimization flow and are designed as a set. Use them together, as you go through the optimization review process.

#### How to Use the Dashboards

The **Consumer \ Update it?** dashboard is a collection of tables (List View) that can be reviewed independently. Click the object name to navigate to the Object Summary page to view more configurations. There can be valid reasons why specific configurations are not followed. It is recommended that you discuss best practices with VMware.

- Outdated vSphere Components Widgets:
  - Lists all the vCenter Servers versions that are not 6.7 or 7.0.
  - Lists all the ESXi host versions that are not 6.5, 6.7, or 7.0.
  - Lists all the vSAN ESXi host versions that are not 6.7 or 7.0. A more stringent filter is applied for vSAN because of a relatively higher maturity in the latest release. From VMware Aria OperationsVMware Cloud Foundation Operations and VMware Cloud Foundation Operations for logs, there are more counters, properties, and events that improve monitoring and troubleshooting.
  - Lists all the vSphere distributed switches, regardless of the version.
  - You should tailor the filter to fit your operational needs.
- Outdated Server BIOS Widget:
  - Lists all the ESXi hosts regardless of the BIOS version. Edit the widget and tailor the filter to fit your operational needs.
- Other than customizing the existing widgets, consider adding the following checks:
  - ESXi hosts with outdated hardware, using a filter based on your environment.
  - ESXi hosts that are no longer on warranty. Create a custom property to capture the end of warranty.
  - Physical storage arrays with outdated firmware, model, and an expiring warranty.
  - Physical network switch with an outdated OS version and hardware model

    **NOTE**
    Install the relevant management pack for the last two points.

#### Points to Note

See the **Points to Note** section as specified in the Consumer \ Correct it? dashboard. This dashboard follows the same design considerations, and as a result, shares limitations and customization ideas.

## Cost Dashboards

The dashboards in the cost category cater to cloud administrators who are responsible for managing the expenses related to your cloud infrastructure. Using Cost dashboards, you can compare the cost of VMware cloud infrastructure with other cloud platforms. You can analyze the cloud comparison results and identify the opportunities to manage your cloud resources efficiently.

### Consumer Layer

The consumer layer dashboards of VMware Aria OperationsVMware Cloud Foundation Operations helps you to know how a customer can do a deeper analysis of the Return on Investment from the consumer perspective.

The available dashboards for consumers are:
- Chargeback VM Price Dashboard

- Showback VM Cost Dashboard
- Showback vSphere Pod Cost Dashboard

### Chargeback (Business Application Price)

In VMware Aria Operations VMs are analyzed financially from the perspective of cost and price.

**Cost**

Cost represents your expenditure to operate the VM for your customer. In VMware Aria Operations, configure Cost Drivers to automatically calculate VM costs based on infrastructure expenses. Cost Drivers encompass server hardware, storage, licenses, applications, maintenance, labor, network, facilities, and additional costs set within VMware Aria Operations.

**Price**

Price is what you charge your customer for using the VM. It can be based on the VM's cost or a predefined rate card, including upcharges, profit, and service charges.

For configuring pricing cards for VMware Aria Automation managed VMs, see the topic, What is Cloud Assembly. For other VMs, the information about configuring pricing cards is in the topic, Add New Pricing Card.

#### How to Use the Dashboard

Select an object in the **Select a Group** widget to view the price of the group. **Price Summary of Selected Group** shows the month to date price of the group. **VM Price Distribution (Top 100)** shows the most expensive VMs in the group. **Powered Off VMs** shows reclaimable VMs and their potential savings. **Idle VMs** shows reclaimable VMs and their potential savings. **VMs with Snapshots** shows reclaimable snapshots and their age. **Price of VMs in the Selected Group** shows the price and configuration of each VM in the selected group.

#### Customizing the Dashboard

The definition of a group is based on several constructs such as Virtual Machine Folder, Custom Groups, VMware Aria Automation Projects, and more. To add more objects for chargeback reporting, edit the **Select a Group** widget and select desired object types under **Output Filter › Basic › Object Types**.

**NOTE**
Objects are not displayed correctly if the widget configuration is not complete. You must complete the object list widget configuration before you can edit the widget.

### Chargeback VM Price Dashboard

The chargeback VM price dashboard lets you know how much you must spend to run a VM on behalf of your customer. In VMware Aria OperationsVMware Cloud Foundation Operations, you can configure the cost drivers and let the system automatically determine how much a VM costs based on your infrastructure requirement. Cost Drivers cover server hardware, storage, licenses, application, maintenance, labor, network, facilities, and additional costs configured within VMware Aria OperationsVMware Cloud Foundation Operations.

Price is what you charge your customer for running their VM. The price of a VM can be based on the cost of the VM or based on a rate card that you define. Prices can include up charges, service charges, and others.

**How to Use the Dashboard**

- Select a Group widget displays the price of the group.
- Price Summary of Selected Group shows the month to date price of the group.
- VM Price Distribution (Top 100) shows the most expensive VMs in the group.
- Powered Off VMs shows reclaimable VMs and their potential savings.
- Idle VMs shows reclaimable VMs and their potential savings.
- VMs with Snapshots shows reclaimable snapshots and their age.
- Price of VMs in the Selected Group shows the price and configuration of each VM in the selected group.

### Showback (Business Application Cost)

This dashboard enables showback of VM cost based on groups.

Improve the accuracy of these costs by editing Cost Drivers. To edit Cost Drivers, see the topic, Editing Cost Drivers. Refer to Overview of Cost Drivers if you are not customizing Cost Drivers.

> **NOTE**
> You must have the Advanced or Enterprise edition license to customize Cost Drivers.

**How to Use the Dashboard**

Select an object in the **Select a Group** widget to view the cost of the group. **Cost Summary (This Month)** shows the month to date cost, potential savings and projected cost of the group. **VM Cost Distribution (Top 100)** shows the most expensive VMs in the group. **Potential Savings (Top 10)** shows the VMs ranked by their potential savings. **Members of the Group (Select to View Trend)** shows the cost and configuration of each VM in the selected group. **Cost Trend of Selected VM** shows the trend of the VMs cost over time.

**How to Customize the Dashboard**

The definition of a group is based on several constructs such as Virtual Machine Folder, Custom Groups, VMware Aria Automation Projects, and more. To add more objects for showback reporting, edit the **Select a Group** widget and select desired object types under **Output Filter › Basic › Object Types**.

> **NOTE**
> Objects are not displayed correctly if the widget configuration is not complete. You must complete the object list widget configuration before you can edit the widget.

### Showback VM Cost Dashboard

The Showback VM Cost dashboard provides a quick Showback of the cost associated with the VMs in a group. Based on the Showback you can improve the accuracy of the costs by editing the cost drivers. Cost drivers that are not customized use reference cost, cost driver customization is available only in Advanced or Enterprise edition of VMware Aria OperationsVMware Cloud Foundation Operations.

**How to Use the Dashboard**

- Select an object in the Select a Group widget to view the cost of the group.
- Cost Summary (This Month) shows the month to date cost, potential savings, and projected cost of the group.
- VM Cost Distribution (Top 100) shows the most expensive VMs in the group.
- Potential Savings (Top 10) shows the VMs ranked by their potential savings.
- Members of the Group (Select to View Trend) shows the cost and configuration of each VM in the selected group.
- Cost Trend of Selected VM shows the trend of the VMs cost over time.

### Showback vSphere Pod Cost Dashboard

The Showback vSphere Pod Cost Dashboard provides a quick Showback of the cost associated with the vSphere Pods in a group. Based on the Showback you can improve the accuracy of the costs by editing the cost drivers. Cost Drivers that are not customized use reference cost, cost driver customization is available only in Advanced or Enterprise edition of VMware Aria OperationsVMware Cloud Foundation Operations.

#### How to Use the Dashboard

- Select an object in the Select a Group widget to view the cost of the group.
- Cost Summary (This Month) shows the month to date cost and projected cost of the group.
- vSphere Pod Cost Distribution (Top 100) shows the most expensive vSphere Pods in the group.
- Idle vSphere Pods shows the vSphere Pods that have been identified as potentially idle.
- Members of the Group (Select to View Trend) shows the cost and configuration of each vSphere Pod in the selected group.
- Cost Trend of Selected vSphere Pod shows the trend of the vSphere Pod's cost over time.

### Provider Layer

The provider layer dashboards of VMware Aria OperationsVMware Cloud Foundation Operations, helps you to know how a customer can analyze the Return on Investment for the virtual infrastructure used in the customer’s environment.

The available dashboards for providers are:

- Assess Cost Dashboard
- Datacenter Cost Drivers Dashboard
- Server Hardware Depreciation Dashboard
- Base Rate Analysis Dashboard
- VM Cost versus Price Dashboard
- Reclaimable Hosts Dashboard

### Assess Cost Dashboards

The **Assess Cost** dashboard provides an overview of the scale of your infrastructure in terms of physical capacity available.

#### Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

#### Widget Information

- You can view the total cost of ownership per month for the infrastructure and the savings opportunities details, if any, for the infrastructure.
- You can view the details of the division of infrastructure investments across all data centers. The dashboard provides the magnitude of each data center in terms of the number of physical servers and virtual machines. It also provides details about the amount of savings that can be achieved from each of these data centers.
- The dashboard displays data about how you invest across clusters of different quality offered across all vCenter Servers.

### Base Rate Analysis Dashboard

The **Base Rate Analysis** dashboard helps you analyze the cost efficiency of your data center.

#### Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

#### Widget Information

- The total cost of ownership is the cost required to run your data center per month. This is derived from the cost drivers.
- The average cost per VM is derived by considering the cost of all the VMs in your environment. The cost of each VM depends on the base rate of the cluster the VM is placed on and its utilization. The base rate of the cluster is computed based on the total cost of ownership and the expected utilization levels of the cluster. Storage base rates are directly obtained from cost drivers.
- If the cluster is running on an allocation-based capacity model, the base rate is derived from the total cost of the cluster and the over-commit ratio. The base rate is indicative of how costly a resource is, on a given cluster.
- A base rate is derived from the total cost and the expected utilization of the cluster.
- A deeper analysis of the base rates can be performed using the CPU, memory, or storage-related widgets, which help rank clusters and datastores relative to their base rates.

### Business Applications Cost vs. Price

The Business Applications showback, chargeback, and cost vs. price dashboard helps analyze the relationship between the cost and pricing of VMs. Use this dashboard to ensure that the chargeback prices for VMs are adequate to cover their operational costs.

#### How to Use the Dashboard

- The **Select a Group** widget lets you select a group of VMs to analyze.
- The **Summary (Month to Date)** widget shows the month to date price and cost.
- The **Members of the Group (Select to View Trend)** shows all VMs in the selected group with their Month to Date Cost, Today's Cost, Month to Date Price, and Today's Price.
- The **Daily Cost and Daily Price** trend chart shows both cost and price over time.

  **NOTE**
  Objects are not displayed correctly if the widget configuration is not complete. You must complete the object list widget configuration before you can edit the widget.

### Datacenter Cost Drivers Dashboard

The **Datacenter Cost Drivers** dashboard provides the cost of different data centers in a private cloud.

#### Customizations Available for Your Use

Certain data centers can be excluded, such as the development data centers that do not have to be expensed, by customizing the views in the widget.

**Widget Information**

- You can select individual data centers to view summary and trends. The summary of the data center costs is grouped into two:
  - Compute. Covers all the costs that are spent on compute related hardware, software, and services.
  - Non-Compute. Covers storage and network.
- Expense trends provide cost variations over a period which indicate infrastructure additions or removal to the data center.
- Cluster expenses indicate the component clusters of a data center that consume the costs. Datastores that represent the storage part of the data center cost are listed alongside.

  **NOTE**
  Network costs are mapped directly to ESXi hosts and hence are costed under compute as well, as of today. This might change in the future.

- When you select a cluster, you can view the component hosts that the cluster is made up of and their monthly depreciated costs. It also provides details on the purchase cost of the server and how many months until it depreciates completely.

  **NOTE**
  Server costs can be suggested out-of-the-box by the system, or can be customized by the user. Depreciation information is not available for servers when the server costs are suggested out-of-the-box by the system. Depreciation information is available for those servers when the server cost is customized by the user.

### Reclaimable Hosts Dashboard

The Reclaimable Hosts dashboard helps you to identify clusters with reclaimable hosts and the potential cost savings from reclaiming the hosts. Reclaimable hosts are identified from the Total Recommended Capacity generated by the AI powered capacity engine in VMware Aria OperationsVMware Cloud Foundation Operations.

**Widget Information**

- The Reclaimable Hosts Cost Pie Chart displays the reclaimable host cost distribution for individual clusters in your virtual environment.
- Potential Savings graph depicts the total cost savings (potential) for all the clusters in your virtual environment for a given period.
- The Top 10 Clusters with Reclaimable Hosts displays the number of reclaimable hosts.
- The Top 10 Clusters with Reclaimable Hosts by Cost displays the top 10 reclaimable hosts by cost.

### Server Hardware Depreciation Dashboard

The Server Hardware Depreciation Dashboard helps you to calculate the depreciation value for server hardware which is marked as owned in Cost Drivers. You can configure the depreciation cost settings as per your business requirement.

**Widget Information**

- The Server Purchase Cost is the total purchase price of all servers as entered in Cost Drivers.
- Accumulated Depreciation is the amount of server purchase costs that have been depreciated according to purchase date and depreciation settings.
- Remaining Depreciation is the amount of server purchase costs left to be depreciated.
- Number of Fully Depreciated Servers identifies servers that have been fully depreciated. These servers may exhibit higher failure rates or have lower capacity. Use What-If scenarios to model the cost and capacity impact of replacing these servers.

### VM Cost vs. Price Dashboard

The VM Cost vs. Price Dashboard helps you to analyze the relationship between cost and price for virtual machines. You can use this dashboard to ensure the price of VMs for chargeback is sufficient to cover the cost of running virtual machines.

**How to Use the Dashboard**

- Select a Group allows selection of a group of VMs to analyze.
- Summary (Month to Date) shows the month to date price and cost.
- Members of the Group (Select to View Trend) shows all VMs in the selected group with their Month to Date Cost, Today's Cost, Month to Date Price, and Today's Price.
- Daily Cost and Daily Price trend chart shows both cost and price over time.

### Cost Optimization

The cost optimization dashboards are as follows:

- Cost Optimization Dashboard
- Potential Cost Savings Dashboard
- Realized Cost Savings Dashboard
- Total Cost of Ownership Dashboard
- VM Rightsizing Details Dashboard

### Cost Optimization Dashboard

The Cost Optimization dashboard helps you to measure the return on investment if you use VMware Aria OperationsVMware Cloud Foundation Operations to manage your virtual infrastructure. You can track the total cost of ownership of the entire environment along with potential savings and realized savings from recommendations provided, the dashboard helps you quantify the cost efficiency and cost savings over time.

Potential Savings is a summary of all cost savings opportunities identified by VMware Aria OperationsVMware Cloud Foundation Operations. Realized Savings is a summary of cost savings from actions performed that are related to recommendations provided by VMware Aria OperationsVMware Cloud Foundation Operations.

**Widget Information**

- Total Cost of Ownership provides details of the monthly cost of server hardware, licenses, maintenance, facilities, labor, network, storage, and additional costs.
- Average Cost per VM is a good indicator of cost efficiency over time. It is natural for the cost per VM to go up when new capacity is added and trend downwards as additional capacity is consumed. The goal is to reduce the average cost per VM over time.
- Realize Savings Breakdown shows the cost of reclaimed resources from the VM identified by VMware Aria OperationsVMware Cloud Foundation Operations.
- Potential Savings covers the cost savings opportunities identified by VMware Aria OperationsVMware Cloud Foundation Operations.

### Potential Cost Savings Dashboard

The Potential Cost Savings Dashboard helps you to measure the cost saving as reported by VMware Aria OperationsVMware Cloud Foundation Operations. You can evaluate the potential savings to track recommendations and

improve cost efficiency over time. The dashboard shows both cost savings and capacity savings for idle VMs, powered off VMs, VM snapshots, orphaned disks, oversized VMs, and reclaimable hosts.

**Widget Information**

- The Cost Savings Breakdown widget displays potential savings and reclaimable capacity for idle VMs, powered off VMs, VM snapshots, orphaned disks, oversized VMs, and reclaimable hosts. You can also view the allocation changes for the oversized VMs.
- Reclaimable widget provides the metric details for the reclaimable vCPU, reclaimable memory, and reclaimable disk space.
- The Optimization Opportunities Breakdown widget covers the projected costs to improve performance as identified by VMware Aria OperationsVMware Cloud Foundation Operations.
- Allocation Changes for Undersized VMs shows the number of vCPUs and GB of memory to add to undersized VMs.

### Realized Cost Savings Dashboard

The Realized Cost Savings Dashboard helps you to quantify the realized cost savings from actions performed that are related to recommendations provided by VMware Aria OperationsVMware Cloud Foundation Operations. You can analyze the realized savings to track improvements to cost efficiency over time. Realized savings covers powered off VMs that were flagged as idle, deleted VMs that were flagged as idle or powered off, deleted snapshots that were flagged as reclaimable, deleted disks that were flagged as orphaned, oversized VMs that were rightsized, and deleted hosts that were flagged as reclaimable.

**Widget Information**

- Realized Savings covers the cost savings from reclamation opportunities recommended by VMware Aria OperationsVMware Cloud Foundation Operations.
- The Reclaimed Capacity shows the amount of capacity that was reclaimed based on recommendations from VMware Aria OperationsVMware Cloud Foundation Operations.
- Allocation Changes for Oversized VMs shows the number of vCPUs and GB of memory removed from formerly oversized VMs.
- Cost of Deleted VMS shows the cost of all deleted VMs in the past 30 days, shows the cost of all deleted VMs (by cluster) for the past 30 days, and shows the year-to-date cost of all deleted VMs.

### Total Cost of Ownership Dashboard

The Total Cost of Ownership dashboard helps you to understand the total cost of ownership of your environment from multiple perspectives. You can use this dashboard to learn how cost drivers, capacity, and data centers affect the total cost of ownership.

**Widget Information**

- Cost Driver Breakdown widget shows how cost drivers affect the total cost of ownership.
- Cost of Capacity Used and Capacity Remaining widgets shows cost breakdown by the cost of capacity used and the cost of capacity remaining.
- Cost per Datacenter widget shows how the costs broken down per data center.

### VM Rightsizing Details Dashboard

The VM Rightsizing Details dashboard provides an overview of the rightsizing recommendations for Undersized VMs and Oversized VMs. Rightsizing is defined as changing the amount of resources allocated to a VM based on the

Recommended Size for a VM. Recommended Size is the maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining.

**How to Use the Dashboard**

- Select a Cluster, Datacenter, or World Object.
- Select an Undersized VM to view the recommendations.
- Select an Oversized VM to view the recommendations.
- Search for a VM to view the recommendations.

## Performance Dashboards

Performance is about ensuring workloads get the necessary resources. Key Performance Indicators (KPI) can be used to identify performance problems related to workloads. Use these KPIs to define SLAs associated with tiers of service. These dashboards use KPIs to display the performance of workloads at the consumer layer and the aggregate performance of workloads at the provider layer.

SLA is the formal business contract that you have with your customers. Typically, SLA is between the IaaS provider (the infrastructure team) and the IaaS customer (the application team or business unit). Formal SLA needs operational transformation, for example, it requires more than technical changes and you might need to look at the contract, price (not cost), process, and people. KPI covers SLA metrics and additional metrics that provide early warning. If you do not have an SLA, then start with Internal KPI. You must understand and profile the actual performance of your IaaS. Use the default settings in VMware Aria OperationsVMware Cloud Foundation Operations if you do not have your own threshold, as those thresholds have been selected to support proactive operations.

The following graphics depict the above relationship.

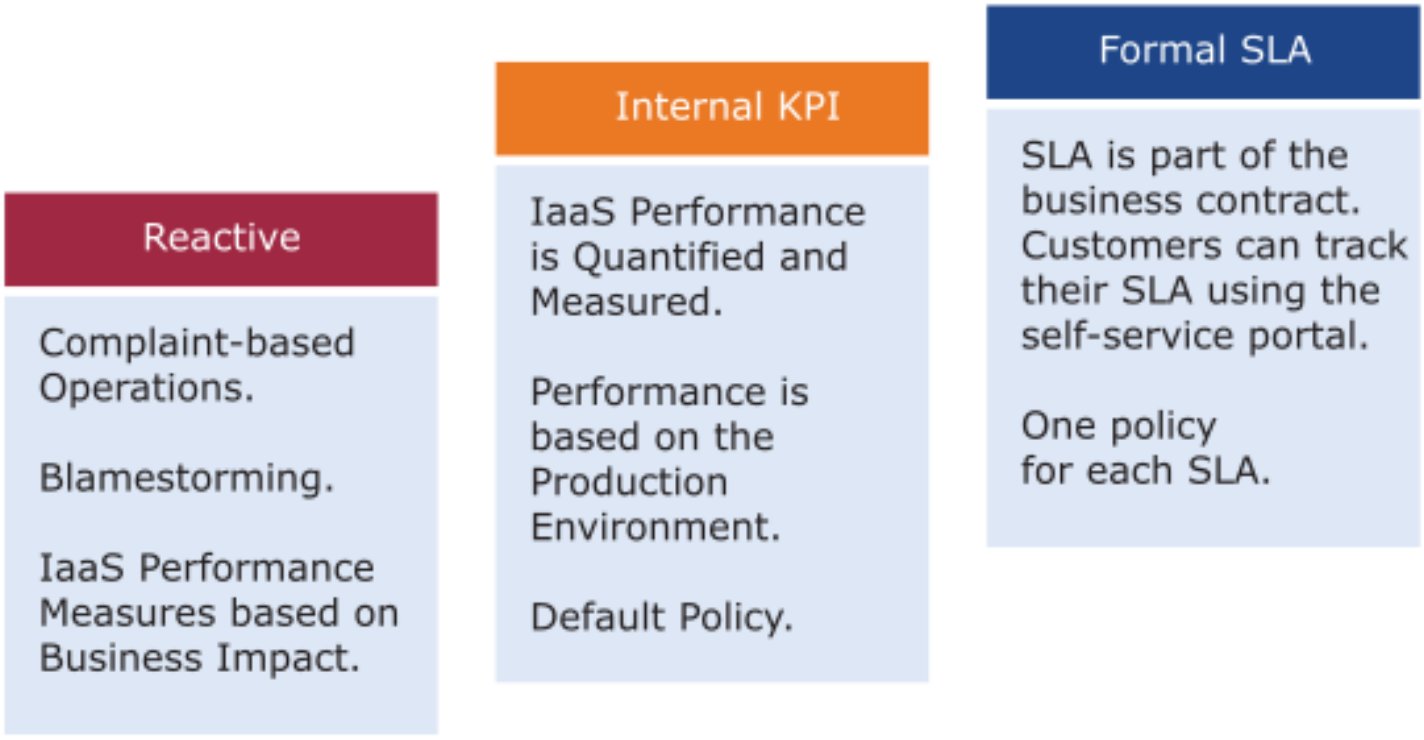


**The Three Processes of Performance Management**

In performance management, there are three distinct processes.

- Planning. Set your performance goals. When you architect a vSAN, you must know how many milliseconds of disk latency you want. 10 milliseconds measured at the VM level (not the vSAN level) is a good start.
- Monitoring. Compare the plan with the actual. Does the reality match what your architecture was supposed to deliver? If not, you must fix it.
- Troubleshooting. When the reality is not according to the plan, you must fix it proactively and not wait for issues and complaints.

To understand what is not healthy for performance management consider the following areas in the given order.

1. Contention: This is the primary indicator.
2. Configuration: Check the version incompatibilities.
3. Availability: Check for soft errors. vMotion stun time, lock up. This requires Log Insight.
4. Utilization: Check this in the end. If the first three parameters are good, you can skip this.

## The Three Layers of Performance Management

There are three main realms of enterprise applications. Each of these realms has its own set of teams. Each team has a set of unique responsibilities and requires the associated skill set. The three realms comprise of Business, Application, and IaaS. Refer to the graphic below to understand the three layers and the typical questions asked on each layer.

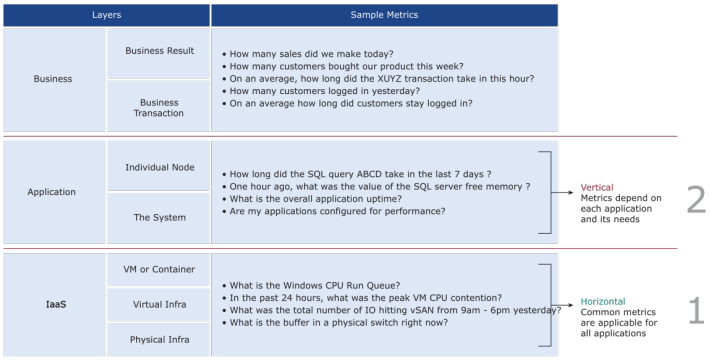


Performance Management is largely an exercise in elimination. The methodology slices each layer and determines if that layer is causing the performance problem. Hence it is imperative to have a single metric to indicate if a particular layer is performing or not. This primary metric is aptly named Key Performance Indicator (KPI).

The upper layer depends on the layer below it, and hence the infrastructure layer is typically the source of contention. As a result, focus on the bottom layer first, as it serves as the foundation for the layer above it. The good part is this layer is typically a horizontal layer, providing a set of generic infrastructure services, regardless of what business applications are running on it.

## The Two Metrics of Performance Management

The primary counter for performance is contention. Most look at utilization, because they fear something wrong might happen if utilization is high. That something is contention. Contention manifests in different forms like, queue, latency, dropped, canceled, and context switch.

However, do not confuse ultra-high utilization indicators as a performance problem. If your ESXi host experiences ballooning, compression, and swapping, it does not mean that your VM has a performance problem. You measure the performance of the host by how well it serves its VMs. While performance is related to the ESXi host utilization, the performance metric is not based on the utilization, instead it is based on contention metrics.

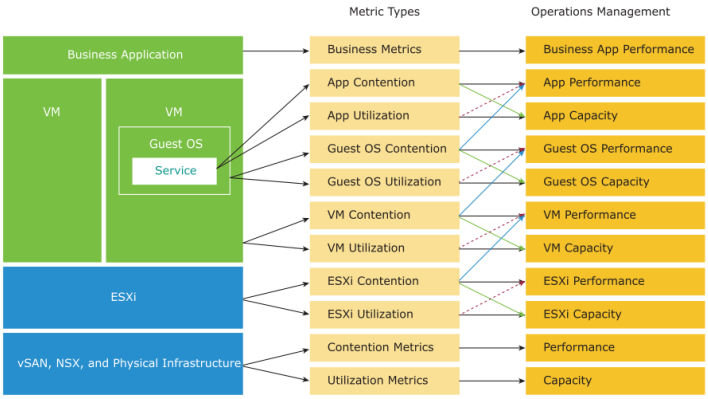


It is possible for VMs in the cluster get affected from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at an individual consumer (VM). The following table shows various possible reasons.

| Infra Configuration | VM and Guest OS Configuration |
|---|---|
| ESXi Settings<br>• Host and BIOS power management causes Frequency to drop.<br>• HT enabled. It looks like twice the capacity, but it is actually 1.25 X throughput.<br>• ESXi - HW compatibility. Driver and firmware are two areas that can impact the performance.<br>• Mismatch of queue depths along the various storage stacks. Must calibrate all the way to the physical array.<br>• vMotion too slow or high stunned time. | VM: Limit, Share, and Reservation<br>• Make sure that no limit is set. CPU ready includes limit.<br>• Make sure that the shares are consistent (as per what the VMs want or you agree to.)<br>• Avoid reservation if possible. This impacts the net available resources for the other VMs. |
| Network<br>• MTU mismatch.<br>• Hops. Especially horse-shoe, or going through multiple ESXi. | Size: NUMA effect. VM spanning NUMA nodes. |
| Cluster Settings<br>• Inconsistent configuration among hosts in a cluster. EVC Mode can play a part if the hosts are from different generations.<br>• Resource Pool<br>– Make sure the shares match the number of VMs.<br>– Make sure that no VM is siblings to RP.<br>• VM- Host Affinity.<br>• DRS Setting. | Snapshot. IO is processes 2x.<br>VM drivers. |
| vSAN<br>• The host where the storage was having performance issues. | Windows or Linux process ping pong, process runaway, and OS level queue. |

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

## VM Performance

Since VM is the most important object in vSphere, it warrants an extra explanation. The graphic below lists the counters you should look at.

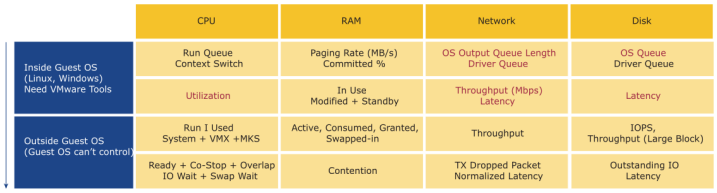


The KPI counters can get technical for some users, so VMware Aria OperationsVMware Cloud Foundation Operations includeS a starting line to get them started. You can adjust the threshold, once you profile your environment. This profiling is a good exercise, as most customers do not have a baseline.The profiling requires an advanced edition.

| | Metric | Green | Yellow | Orange | Red |
|---|---|---|---|---|---|
| Guest OS Contention | Total CPU Run Queue | 0-5 | > 5 | > 10 | > 20 |
| | CPU Context Switch Rate | 0 - 5K | < 25K | < 100K | > 100K |
| | Total Disk Queue Length | 0 - 25 | > 25 | > 50 | > 100 |
| Guest OS Usage | RAM Free (MB) | > 512 MB | > 256 | > 128 | ≤ 128 |
| | RAM Page-in Rate (KB/s) | 0 - 25K | > 25 K | > 50 K | > 100K |
| VM Contention | CPU Co-Stop (%) | 0 - 2.5% | > 1 | > 3 | > 5 |
| | [SLA] CPU Ready (%) | 0 - 2.5% | > 2.5 | > 5 | > 7.5 |
| | Total CPU Overlap (ms) at VM level | 0 - 1000 | > 1000 | > 2500 | > 5000 |
| | CPU IO Wait | 0 - 1000 | > 1000 | > 2500 | > 5000 |
| | [SLA] RAM Contention (%) | 0 - 1% | > 1 | > 2 | > 4 |
| | [SLA] Disk Latency (ms) | 0 - 10 ms | > 10 | > 20 | > 40 |
| | [SLA] Network TX Dropped Packet | 0 | > 0 | > 1 | > 2 |
| VM Usage | CPU Usage (%) | 0 - 85% | > 85 | > 90 | > 95 |

## Performance Metrics

VMware Cloud Foundation OperationsVMware Aria Operations uses the following threshold for internal KPI.

| IaaS | VM Counter | Threshold |
|---|---|---|
| CPU | Ready | 2.5% |
| RAM | Contention | 1% |
| Disk | Latency | 10 ms |
| Network | TX Dropped Packet | 0 |

The table is an example of a stringent threshold. A high standard for performance is used because it is an internal KPI for the consumption of the infrastructure team. It is not an external formal SLA that is confirmed with the customers. There must be a buffer between the internal KPI and the external SLA so that the operations team receive early warnings and has the time to react before the external SLA is breached. A high standard also works from the mission critical point to view to the development environment. If the standard is set to the least performing environment, then it cannot be applied to the more critical development.

A single threshold is used to keep the operations simple. This means that the performance in production is expected to have a higher score than the development environment. The development environment performance is expected to be worse than the production environment, while everything else is equal. A single threshold helps to explain the difference in Quality of Service (QoS) provided by a different class of service. For example, if you pay less, you get a poor performance and if you pay half the price, expect to get half the performance.

The four elements of IaaS (CPU, RAM, Disk, and Network) as mentioned in the table, are evaluated on every collection cycle. The collection time is set at five minutes as it is an appropriate balance for monitoring. If SLA is based on one minute, it is too close and results in either cost increase or reduction in threshold.

**Design Considerations**

All the performance dashboards share the same design principles. They are intentionally designed to be similar, as it is confusing if each dashboard looks different from one another, considering they have the same objective.

The dashboards are designed with separate two sections: summary and detail.

- The summary section is typically placed at the top of the dashboard to provide the overall picture.
- The detail section is placed below the summary section. It lets you drill down into a specific object. For example, you can get the detailed performance report of any specific VM.

In the detail section, use the quick context switch to check the performance of multiple objects during performance troubleshooting. For example, if you are looking at the VM performance, you can view the VM-specific information and the KPIs without changing screens. You can move from one VM to another and view the details without opening multiple windows.

The dashboard uses progressive disclosure to minimize information overload and ensure the webpage loads fast. Also, if your browser session remains, the interface remembers your last selections.

Many of the performance and capacity dashboards share a similar layout since there is a shared commonality between these pillars of operations.

### Guest OS Performance Profiling Dashboard

Use the **Guest OS Performance Profiling** dashboard to know the actual performance of your environment.

Some counters directly impact the performance of Windows or Linux, the operating systems running inside the VM. These KPIs are outside the control of the hypervisor.

Modern operating systems such as Linux and Windows use memory as cache, since it is faster than a disk. Some counters directly impact the performance of Windows or Linux. These KPIs are outside the control of a hypervisor, which means that the ESXi VMkernel cannot control the increase or decrease of the KPI values. The KPI visibility also requires an agent, such as VMware Tools. As a result, they are typically excluded in performance monitoring.

Since they are closer to the applications, it is critical to know their values and establish an acceptable range. The acceptable level of these KPIs among all the VMs in your environment varies. By profiling the actual performance across time and from all VMs, you can establish a threshold that is supported by facts. Since there are 8766 instances of 5 minutes in a month, profiling 1000 VM over a month means you are analyzing 8.8 million datapoints.

**Design Considerations**

The dashboard uses progressive disclosure to minimize information overload and ensures that the webpage loads fast.

In a large environment, loading thousands of VMs increases the loading time of VMware Cloud Foundation OperationsVMware Aria Operations. As a result, the VM is grouped by data center. For a small environment, vSphere World is provided so you can see all the VMs in the environment.

**How to Use the Dashboard**

Select data center from the data centers list. The three tables listing CPU, memory, and disk will show the VMs in the selected data center or vSphere world. Each table shows the highest value in the last one week (2016 datapoints based on five minutes collection cycles), and hence uses the term max as a prefix, for example Max Page-Out/sec or Max Guest OS Disk Queue.

Select any of the VMs in any of the tables. The three line charts are displayed. They are showing data from the same VM to facilitate correlation.

- CPU table widget:

  - The Max CPU Queue column shows the highest number of processes in the queue during the given period. As a best practice, keep the queue below three for each queue. A VM with eight CPUs has eight queues, hence keep this number below 24.
  - The CPU Hyperthreading gives twice the queue as it should as both threads are interspersed in the core pipeline.
  - CPU Context Switch. There is a cost associated with the context switch. There is no guidance for this number, and it varies widely.
- Memory list widget:
  - In memory paging, the modern operating systems (Linux and Windows) use memory as cache, it is much faster than a disk. It proactively pre-fetches pages and anticipates future needs (Windows calls this Superfetch). The rate pages that are being brought in and out can reveal memory performance abnormalities. A sudden change, or one that has sustained over time, can indicate page faults. Page faults indicate that pages are not readily available and must be brought in. If a page fault occurs too frequently, it can impact application performance. While there is no concrete guidance, as it varies by application, you can view a relative size. operating systems typically use 4 KB or 2 MB page sizes.
- Disk list widget:
  - Disk queues are queued IO commands that are not sent to the VM. They have been retained inside the Guest OS (either at a kernel level or a driver level). A high disk queue in the guest OS, accompanied by low IOPS at the VM, can indicate that the IO commands are stuck waiting on processing by Windows/Linux. There is no concrete guidance regarding these IO commands threshold as it varies for different applications. You should view this with the Outstanding Disk IO at the VM layer.

**Points to Note**

- These Guest OS widgets do not appear unless the vSphere pre-requisites are met. For more information, see KB article 55697.
- Once you determine an acceptable threshold for your environment, consider adding thresholds to the table so you can easily view the VMs that exceed a threshold.
- The CPU queue is the sum from all virtual CPUs. A larger VM can tolerate a higher queue as it has more processors. If you want to compare VMs of different sizes, create a super metric that calculates the queue per vCPU. For more information, see Create a Super Metric.
- Group the VM by clusters of the same class (for example, Gold), so you can see the profile for each environment.
- For a smaller environment, consider changing the table from listing data centers to listing clusters.

### Network Top Talkers Dashboard

Use the **Network Top Talkers** dashboard to monitor network demand in your IaaS. In a shared environment, a few VMs generating excessive activity can impact the entire data center. While a single VM might not cause a serious problem, a few of them can.

**Design Considerations**

The **Network Top Talkers** dashboard helps you analyze how hard these VMs hit your IaaS. It classifies the workload into two: short bursts and sustained hits. A short burst lasts for a short period, maybe for a few minutes. A sustained hit can last for an hour and cause serious problems.

The **Network Top Talker** dashboard forms a pair with the **Storage Heavy Hitter** dashboard. To understand the IO demand in your environment, use both of them concurrently.

The **Network Top Talkers** dashboard displays sustained hits that last for an hour, as they can cause serious problems in a shared IaaS environment. You can identify the villain VMs and compare their demands with the capabilities of the underlying IaaS.

### How to Use the Dashboard

The dashboard shows the current workload. This is the total network load (received and transmitted) from all the vSphere environments monitored by VMware Aria OperationsVMware Cloud Foundation Operations. The idea is to give you an indicator on how hard the overall load is.

- Select a data center from the data centers list.
  - The columns show the number of clusters, ESXi hosts, and VMs for each data center. The VM count includes the powered off VM. To only see the running VM count, edit the widget.
  - If you want to see information from all the data centers, select the vSphere world row.
  - Upon selection, the Total Demand Line chart and the Top Talkers tables fill up.
- Total Demand Line Chart
  - The total throughput (received and transmitted) in the selected data center.
  - Displays both, the five minute peak and the hourly average in one line chart. You can click the metric name to hide it.
- Top Talkers Table
  - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important. For example, an ESXi with 2 x10 GB port can theoretically handle 20 GB TX + 20 GB RX as its full duplex.

### Points to Note

- Understanding high demand helps you monitor IaaS and plan your capacity. IaaS provides four services, CPU, memory, disk, and network. While CPU, memory, and disk are bound, an active VM can consume all your network bandwidth, packet per second capacity, and the storage IOPS capcity. A VM with 4 vCPU and 16 GB memory cannot consume more than this amount, the same applies to disk space. A VM configured with 100 GB disk space cannot consume more than that.
- Network throughput, disk throughput, and disk IOPS can spike as their physical limits are very high per VM. This means that IaaS has enough capacity for all workloads and performs well until the VMs start consuming abnormally high amounts of network and disk bandwidth.

## Storage Heavy Hitters Dashboard

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard. To understand the IO demands in your environment, use both of them together. If you are using ethernet-based storage, storage traffic runs over the same physical network as your ethernet-based network traffic.

### Design Considerations

The **Storage Heavy Hitters** dashboard forms a pair with the **Network Top Talkers** dashboard, so they share a consideration behind their design. For more information, see Network Top Talkers Dashboard.

### How to Use the Dashboard

- See the **Network Top Talkers** dashboard as they have the same design.
  - The main difference between **Storage Heavy Hitters** and **Network Top Talkers** is that the storage IO has two dimensions: IOPS and throughput.
  - Network IO does not have the IOPS dimension as the packet size is identical (1500 bytes being the standard packet, and 9000 bytes being the jumbo frames).
  - Storage IOPs and throughput are related, so use both to gain insight, they should display a similar pattern. If not, that indicates varying block sizes. For example, a throughput spike without an accompanying IOPs spike indicates large block sizes.
- Which VMs hit the storage the hardest.
  - The table shows the most demanding VM. You can identify the villain VM and compare their demands with the capabilities of the underlying IaaS. Knowing the infrastructure capability is important, because different classes of SSD have different IOPS and throughput capabilities.

After identifying the villain VM, talk to the VM owners if the numbers are excessive during peak hours and identify the reasons behind the excessive usage. You must ensure that they do not create a hot spot. For example, vSAN cluster with > 100 disk can handle numerous IOPS but if the VM objects are only on a few disks, those disks can become a hot spot.

**Points to Note**

- Interpreting IOPs and throughput metrics depends on your underlying physical storage. For visibility into this hardware layer, add physical storage metrics to the dashboard.

### VM Contention Dashboard

The **VM Contention** dashboard is the primary dashboard for VM performance. It is designed for VMware administrators or architects. It can be used for both, monitoring, and troubleshooting. Once you determine that there is a performance issue, use the **VM Utilization** dashboard to see if the contention is caused by high utilization.

**Design Considerations**

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

For understanding the performance concept of the selected counters and their thresholds, see the Performance Dashboards

**How to Use the Dashboard**

- Select a data center from the data center table.
  - For a smaller environment, select vSphere World to see all the VMs from all the data centers.

    **NOTE**
    The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.

- The two bar charts are automatically shown.
  - Use them together to get an insight about your CPU readiness and your Memory contention analysis. Analyze how the cluster serves the VMs. For each VM, it picks the worst metric in the last 24 hours. By default, VMware Aria OperationsVMware Cloud Foundation Operations collects data every 5 minutes, so this is the highest value among 288 datapoints. Once it has the value from each VM, the bar charts puts each VM in the respective performance buckets. The threshold in the buckets considers best practices, hence they are color coded.
  - For any critical environment, expect that all the VMs are served well by the IaaS. You must see green on both distribution charts. For development purposes, you can tolerate a small amount of contention in both CPU and Memory.
- VM Performance in selected Data Center.
  - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.
  - The table is sorted by KPI Breach columns, directing your attention to the VMs that are not served well by the IaaS.
  - The table shows the hostnames known by Windows or Linux. This is the name that the application team or VM owner knows, as they might not be familiar with the VM name.
  - The rest of the columns show performance counters. Because the goal is proactive monitoring, the counters are the worst and not the average, during the monitoring period. Because the operations context here is

performance, not capacity, the table considers the last 24 hours only. Daily use is encouraged as any activity older than 24 hours is considered irrelevant from a performance troubleshooting viewpoint.
  - The column KPI Breach counts the number of SLA breaches in any given 5 minutes. As a VM consumes four resources of IaaS (CPU, memory, disk, and network), the counter varies from 0–4, with 0 being the ideal. The value 4 indicates that all 4 IaaS services are not delivered. The same threshold is used regardless of class of service, as this is an internal KPI, not an external SLA. Your internal threshold should be more stringent, so that you have a reaction time.
- Select a VM from the table.
  - All the health charts show the KPI of that VM.
  - The health charts display the last value, lowest value, and the peak value. Expect that the peak is within your threshold.

**Points to Note**

- This dashboard uses Guest OS counters and VM counters appropriately. The two layers are distinct layers, and they each provide a unique visibility that the other layers might not give. For example, when the VMkernel de-schedules a VM as it has to process something else (for example, other VM, kernel interrupt). The Guest OS does not know the reason. In fact, it experiences frozen time for that particular vCPU running on the physical core and experiences time jumps when it is scheduled again.
- Guest OS counters logically require VMware Tools.
- The health chart is color coded. Change the settings if it does not suit your environment. If you are unsure of what suitable numbers to set for your environment, profile the metrics. The Guest OS Performance Profiling dashboard provides an example of how to profile metrics.
- For a smaller environment with one or two data centers, change the filter from data center to cluster. Once you are list a cluster, you can then add the cluster performance (%) metric and sort them in an ascending order. This way the cluster that needs immediate attention is on the top.
- If you have a screen real estate, group the VMs by cluster or by ESXi host. This way, you can quickly see if the problem is in a particular cluster or ESXi host.
- Change the default timeline from one week to one day as and when required to suit your operations.
- If you navigate a lot to the **VM Utilization** dashboard from this dashboard, add a connection using the dashboard to dashboard navigation feature. For more details, see Dashboard Navigation Details.

### VM Performance Dashboard

Use the **VM Performance** dashboard to find out if a VM has a performance problem. As a first step, when a VM has a problem, verify if other VMs have the same problem. If the problem is widespread the root cause is not with the VM.

**How to Use the Dashboard**

The **VM Performance** dashboard is organized into sections for ease of use.

- Select a data center from the **Datacenters** widget. To find out if there is a performance problem, what the problem is, and the extent of the problem, use the following three bar charts together: **Are VMs facing CPU Ready**, **Are VMs facing Memory Contention**, **Are VMs facing Disk Latency**. Each bar chart analyzes how the VMs are served by the cluster. These bar charts indicate if the VMs are waiting for CPU resources, facing memory contention, or disk latency. For each VM, it picks the worst metric in the last 24 hours. By default, VMware Aria OperationsVMware Cloud Foundation Operations collects data every 5 minutes, so this is the highest value among 288 datapoints (12 x 24 = 288). Once it has the value from each VM, the bar charts put each VM in the respective performance buckets. The threshold in the buckets considers best practices, and hence they are color coded. For each bar chart, you can change the time period to the period of your interest. The maximum number is then displayed. The value is the worst 20-seconds, within the 5-minute collection time period. For your mission-critical environment, you must expect that all the VMs are being served well by the IaaS. If you see green on the distribution charts, you do not have to analyze further.

  For development, you may tolerate a small amount of contention in both CPU and Memory as you need to balance cost.

You can also change the filter from data center to cluster. If you are listing clusters, you can then add the cluster performance (%) metric and sort them in ascending order. This way the cluster that needs immediate attention is on top.

You can click on the bar to see the list of VMs under that performance bucket. From there, you can select a VM, and it's KPI is automatically displayed on the lower section of the dashboard.

- **Multiple VM Analysis**
  When you select a data center from the **VMs Performance in selected Datacenter** widget, the table listing all the VMs in the data center is displayed.

  The table is sorted by the KPI Breached column, directing your attention to the VMs that are not served well by the IaaS. The column counts the number of SLA breaches in any given 5-minute period. It is based on the counter `Performance \ Number of KPIs Breached`. As a VM consumes four resources of IaaS (CPU, memory, disk, and network), the counter varies from zero through four, with zero being the ideal. The value four indicates that all four IaaS services are not delivered. The same threshold is used regardless of class of service, as this is an internal KPI, and not an external SLA.

  Because the goal is proactive monitoring, as opposed to reactive troubleshooting, the counters show the worst value instead of the average of the monitoring period.
- **Per VM Analysis**
  When you select a VM from the table, the CPU, memory, disk, and network performance charts are automatically displayed, each widget showing the KPIs of that VM.
- **Alerts**
  The relevant alerts are displayed automatically. You can view the settings by editing the widget, and adjust them accordingly to fit your operational needs.
- **Virtual Disks**
  A VM can have many disks, and it is possible that these disks may have different performance levels. The table lists the individual virtual disks and their contention and utilization metrics.
- **Configuration**
  The relevant configuration of the selected VM is displayed. You can customize as appropriate.
- **Relationship**
  From the VM, you can navigate to the parent cluster or datastore. Use the **Relationship** widget to navigate and auto select the associated cluster or datastore.

### VM Utilization Dashboard

The VMware administrator uses the **VM Utilization** dashboard with the **VM Contention** dashboard for managing performance.

#### Design Considerations

Use the **VM Utilization** dashboard to identify virtual machines with a high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when a queue develops inside the Windows or Linux operating systems. By default, VMware Cloud Foundation OperationsVMware Aria Operations has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

#### How to Use the Dashboard

- Select a data center from the data center table.
  - For a smaller environment, select vSphere World to see all the VMs from all the data centers.

**NOTE**
The count of VMs includes the powered off VMs too. To exclude powered off VMs, modify the widget and select the running VM metric.

- VM Peak CPU Usage (%).
  - There is no peak memory use as it is not applicable. Memory is a form of storage, for example consider a hard disk occupied space. A 90% utilization of the total space is not slower than 10%. This means that the issue is related to capacity issue and not performance.
  - The bar chart is color coded using five colors instead of four. The color gray is introduced to convey any wastage. Resources that are hardly utilized do not signify that the performance is at its peak. It can also mean the opposite. For example, if a VM needs 1+ vCPU, configuring it with 2 CPUs results in better performance instead of configuring it with 128 CPUs.
- VM Peak Utilization.
  - Analyze by data center as performance problems tend to be isolated in a single physical environment. For example, a performance problem in country A typically does not cause a performance problem in country B.
  - The table focuses on peak utilization, because the context is performance and not capacity.
- Select a VM from the table.
  - All the health charts show the KPI of that VM.
  - Compliment the free memory with the memory IOPS or the memory throughput metric. The metrics in a gigabyte measure the space, and not the speed. Memory is a form of storage, so what you must measure is the rate, for example, read-write per second.

**Points to Note**

- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the VM Contention Dashboard.

### Troubleshoot an Application Dashboard

The VMware Aria Operations Application Management Pack provides discovered applications to be managed in VMware Aria OperationsVMware Cloud Foundation Operations. Using the **Troubleshoot an Application** dashboard, users can see the applications and the relevant metrics and alerts for the selected application. The dashboard also displays its relationship to the infrastructure. In the list of metrics, select a metric to see its trend over time.

### Cluster Contention Dashboard

The **Cluster Contention** dashboard is the primary dashboard for vSphere cluster performance. It is designed for VMware administrators or architects. It can be used for both, monitoring and troubleshooting. Once you determine that there is a performance issue, use the **Cluster Utilization** dashboard to see if the contention is caused by high utilization.

**Design Considerations**

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed for daily use, hence the views are set to show data for the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

Utilization of the cluster is not shown in the **Cluster Contention** dashboard. You must separate the two concepts: utilization and contention. Performance and capacity are different concepts managed by two separate teams. Both CPU and memory are also shown separately. You can have a problem with one, without any issue in the other. CPU is more common as memory tends to have a lower overcommit ratio.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

**How to Use the Dashboard**

- Average Cluster Performance (%).
  - This is the primary KPI for your entire IaaS. It plots how your IaaS is performing every 5 minutes, giving you the trend view of the overall performance.
  - The metric itself is simply the average of the Cluster KPI / Performance (%) metric. This performance metric in turn averages the VM Performance / Number of KPIs Breached metric from all the running VMs in the cluster. Hence a value of 100% indicates that every running VM in the cluster is served well.
  - As this KPI takes into account every running VM in your environment, the number should be steady. The analogy in real life is the stock market index. While individual stocks can be volatile, overall the index should be relatively steady on a 5 minutes by 5 minutes basis.
  - The relative movement of the metric is as important as the absolute value of the metric. Your absolute number might not be as high you want it to be, but if there are no complaints for a long time, then there is no urgent business justification to improve it.
- Clusters Performance.
  - It lists all the clusters, sorted by the least performing cluster in the last one week. You can change this time period.
  - The worst performance shows the lowest number in the time period. As VMware Aria OperationsVMware Cloud Foundation Operations collects data every 5 minutes, there are 12 x 24 x 7 = 2016 data points in a week. This column shows the worst point among these 2016 datapoints.
  - A single number among 2016 datapoints can be an outlier that needs to be complemented with another number sometimes. A logical choice is the average of these numbers. For the average performance to be low, a lot of criterias have to be low. Waiting for the average causes a delay in your operations, and rise in complaints. For performance monitoring, the 95th percentile is a better summary than the average.
  - Your cluster should function at a 100% and perform its fuctions as planned.
- Select a cluster from the table.
  - All the health charts show the KPI of the selected cluster.
  - For performance, it is important to show both the depth and breadth of the performance problems. A problem that impacts one or two VMs requires a different troubleshooting than a problem that impacts all the VMs in the cluster.
  - The depth is shown by reporting the worst among any VM counter. So the highest value of VM CPU Ready, VM Memory contention, and VM Disk Latency among all the running VMs are shown. If the worst number is good, then you do not need to look at the rest of the VMs.
  - A large cluster with thousands of VMs can have a single VM experiencing poor performance while 99.9% of the VM population is fine. The depth counter might not report that most VMs are fine. It only reports the worst. This is where the breadth counters come in.
  - The breadth counters report the percentage of the VM population that is experiencing performance problem. The threshold is set to be stringent, as the goal is to provide early warning and activate proactive operations.

**Points to Note**

It is possible for VMs in the cluster to suffer from poor performance, while the cluster utilization is low. One main reason is cluster utilization looks at the provider layer (ESXi), while performance looks at individual consumer (VM). The following

table shows various possible reasons.

| Event | Aware? |
| --- | --- |
| Power Management | No |
| HT | No |
| Ready | No |
| Co-Stop | No |
| System | No |
| Steal | No |
| IO Wait | No |
| Memory Wait | No |

From the performance management point of view, the vSphere cluster is the smallest logical building block of the resources. While the resource pool and VM Host affinity can provide a smaller slice, they are operationally complex, and they cannot deliver the promised quality of IaaS service. Resource pool cannot provide a differentiated class of service. For example, your SLA states that gold is two times faster than silver because it is charged at 200% more. The resource pool can give gold two times more shares. Whether those extra shares translate into half the CPU readiness cannot be determined up front.

Certain settings such as DRS automation level and the presence of many resource pools can impact performance. Consider adding a property widget to show the relevant property of a selected cluster, and a relationship widget to show resource pools.

For a large environment with many clusters, add a grouping to make the list more manageable. Group it by class of service, so you can focus more on the critical clusters.

**Cluster Performance Dashboard**

The **Cluster Performance** dashboard combines the functionality of the **Cluster Contention**, **Cluster Utilization**, **ESXi Contention**, and **ESXi Utilization** dashboards.

**How to Use the Dashboard**

The **Cluster Performance** dashboard is organized into sections for ease of use.

- **Overall Analysis**
  The **Average Cluster Performance** health chart is green when all the clusters are performing well. If the clusters are unable to serve the VMs well, all the clusters are no longer green, with a few occurrences of red.

  As the chart displays all the clusters, it uses the vSphere World object. This object is the parent of the vCenter object, and so it displays all the clusters from all the vCenter Servers.

  The metric used is `Performance \ Clusters Performance (%)` and is the primary KPI for your entire IaaS. It plots how your IaaS is performing every five minutes, giving you the trend view of overall performance.
- **Multi-Cloud Analysis**
  If the health chart is not green, and you want to find out which clusters are not performing, use the **vSphere Clusters** widget. The table lists all the clusters, starting with the cluster with the lowest performance. By default, the data displayed is from the last 24 hours. The **Worst Performance** column displays the lowest number in the time period. By default, VMware Aria OperationsVMware Cloud Foundation Operations collects data every 5 minutes, so this is the lowest point among 288 datapoints (12 x 24 = 288).

  The **Worst Performance** column displays the lowest performance in the last period, specified under **Time Settings**.
- **Per-Cluster Analysis**
  Select a cluster from the **vSphere Clusters** widget to see the trend over time. After you determine the cluster you want to investigate, review the five scoreboards: CPU, memory, disk, network, and others.
- **VM Shares**
  A common root cause for uneven performance problems is uneven shares. Each slice in the pie chart must correspond to a class of service. If the entire cluster is serving one class, then you should see a simple circle with no slices.
- **Resource Pool Analysis**
  Resource pool is another common reason behind uneven VM performances. A cluster with too many resource pools makes performance management difficult.

  The **Resource Pools in the cluster** provides a table listing all the resource pools.
- **ESXi Analysis**
  A cluster is a collection of ESXi hosts and the performance can be affected by uneven performance among the member hosts. You can drill down from a cluster to the ESXi hosts. The **ESXi Hosts in the Selected Cluster** widget lists all the ESXi hosts in the cluster, sorted by the worst performance in the last 24 hours. If the table displays values in green, you do not have to analyze further.

  You can change the time period to the period of your interest. The maximum number will be reflected accordingly.

  The table helps you quickly compare the performance of each ESXi. You can also see the performance over time, to see a trend.

Certain settings such as power management and hyper threading can impact performance. The **ESXi Hosts in the Selected Cluster** widget displays the relevant property of a selected ESXi Host.

- **VM Analysis**
  When you select a cluster or ESXi from the **Running VMs in the selected Cluster or ESXi widget** the VMs that are running are automatically listed. Use this table to verify if the cluster or host performance problems were caused by VM configuration and usage. It is possible that the VM was not on the same host at the time of problem, due to vMotion.
  To drill down into a particular VM, select it and click the double arrow before the widget title.
- **Datastore Analysis**
  Use the **Shared datastores in the Cluster** widget to see a list of shared datastores accessible by hosts in the cluster. You can also drill down to the selected datastore.

### Cluster Utilization Dashboard

The VMware administrator uses the **Cluster Utilization** dashboard with the **Cluster Contention** dashboard for performance management.

#### Design Considerations

This dashboard supports the **Cluster Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted especially when VMs experience a contention. By default, VMware Cloud Foundation OperationsVMware Aria Operations has a 5-minutes collection interval. For five minutes, there may be 300 seconds worth of data points. If a spike is experienced for a few seconds, it may not be visible if the remaining of the 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

#### How to Use the Dashboard

- CPU(%) and Memory (%).
  - Review the CPU and Memory distribution charts for an overview of the CPU and memory utilization of the clusters.
  - The highest metric in the last one week is used. Average or 95th percentile is not used as this is utilization and not contention. High utilization does not mean bad performance.
  - One week is used instead of one day to give you a longer time horizon and covers the weekend. Adjust the timeline as you deem fit for your operations.
  - Expect memory to be higher than CPU, as it is a form of cache. The Memory Consumed counter is used as it is more appropriate than the Memory Active counter.
  - Low utilization can actually indicate bad performance, as not much of real work gets done. The chart uses the dark gray color for low utilization.
- Clusters Utilization.
  - The cluster utilization table lists all the clusters, sorted by the highest utilization in the last one week. If the table displays the green color, then there is no need to analyze further.
  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select a cluster from the table.
  - All the utilization charts show the key utilization metrics of the selected cluster.
  - For memory, the high utilization counters are explicitly shown, Balloon, Compressed, and Swapped. Notice they exist even though utilization is not even at 90%, indicating high pressure in the past. If you look only at utilization, you might think you are safe.
  - The line charts show both average and highest among ESXi hosts in the cluster. The reason is unbalanced and it is not rare. There are many settings that can contribute to it (for example, DRS settings, VM Reservation, VM – Host Affinity, Resource Pool, Stretched Cluster, and Large VMs).
  - The disk IOPS is split into read and write to gain insight into the behavior. Some workload is read oriented, while others are write oriented.

  - The disk throughput is not shown as it sums all the traffic. In reality, each ESXi host has its own limit.
  - The vMotion line chart is added, as a high number of vMotion can indicate that the cluster load is volatile, assuming the DRS Automation level is not set to the most sensitive setting.

**Points to Note**

- If your operations team have some forms of standardization that utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold.
- Consider adding a third distribution chart. Show the balloon counter in this third chart, as it complements the consumed counter. If there is no ballooning, a high consumed value is in fact better than a lower value.
- The workload metric can exceed a 100% because it is demand / usable capacity * 100. This can happen if you have four hosts in a cluster with each host running at 100% demand and admission control is set to 50%.
- The **VM Utilization** dashboard complements the **VM Contention** dashboard. For more information, see the points to note in the Cluster Contention Dashboard.

### VM Rightsizing Dashboard

The **VM Rightsizing** dashboard helps you adjust the VM size for optimal performance at the lowest cost. It covers both undersized and oversized scenarios. The dashboard looks at the long term trend and covers both undersized and oversized scenarios. This dashboard is designed for the Capacity and the Operations teams, as rightsizing a VM helps in the day-to-day performance.

**How to Use the Dashboard**

- **Overall Analysis**
  The scoreboard provides a summary of the total undersized and oversized CPU and memory.

  You can select either a data center or a cluster. In most cases, rightsizing analysis should be done at the cluster level as VMs typically do not move inter-cluster. The counters are displayed to provide better context. Focus on reclaiming VM capacity in a cluster that is low on capacity remaining.

  The distribution charts that display rightsizing are automatically displayed. Other than the bar charts, the **Undersized VM** and **Oversized VM** tables list the actual VMs.
- **VM Analysis**
  Select a VM to investigate further. The utilization is automatically displayed. VM usage reflects the amount of capacity consumed. This is based on the aggregate vCPU usage at five minute granularity, which provides a clear understanding of capacity used.

  Rightsizing a VM can help improve performance for the VM as well as the cluster. Because of this, metrics that show performance bottlenecks, such as CPU Ready and CPU Run Queue, are provided to help you confirm how rightsizing may result in less contention and better performance.

  Memory utilization is collected from the guest OS via VMTools. If guest OS metrics are not available, then the memory configured value is used instead. Rightsizing memory improves performance by reducing memory ballooning and contention. For example, VMs with over-provisioned memory are more likely to experience ballooning.

### Datastore Performance Dashboard

Use the **Datastore Performance** dashboard to view performance problems related to storage such as high latency, high outstanding IO, and low utilization. This dashboard is designed for both the VMware administrator and the Storage administrator, to foster a closer collaboration between the two teams. Local datastores are treated separately.

### Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

### How to Use the Dashboard

The **Datastore Capacity** dashboard is layered, gradually providing details as you work top-down in the dashboard.

- **Overall Analysis**
  Select a data center from the **Datacenters** table. The three bar charts which are, **VM Performance**, **Read Performance**, and **Write Performance** provide an overall analysis of the datastore performance in a given vCenter data center or vSphere World. They work together to provide better insight. Just like other performance charts, the value displayed is the worst value during the time period. After you select a data center, if the Worst VM Disk Latency displays `No data to display`, it means that you have no observed latency issues for VM disk performance.

  The **VM Performance** chart displays the kind of latency and how many VMs experience that kind of latency. The **VM Performance** chart is your primary chart as it measures latency at the VM level. The **Read Performance** and **Write Performance** charts measure latency at the datastore level, which means they are the normalized average of all VMs in that datastore. Expect the **VM Performance** chart to be higher than the **Read Performance** and **Write Performance** charts. Read and write latency are displayed separately for better insight.
- **Datastore Analysis**
  The **Datastores Performance** table automatically lists all the shared datastores in the data center or vSphere World. Both the worst (peak) performance and the 95th percentile are displayed. If the latter is close to the peak and it is also high, then it is a sustained problem. If the latter is low, then the problem is for a short duration. The table is color coded. Select a datastore that you want to troubleshoot. The relevant metrics and configuration are displayed.

- **VM Analysis**
  The list of VMs running in the selected datastore is displayed, with the relevant contention and utilization counters. Select the VM that you want to troubleshoot. The contention and utilization of the VM are automatically displayed. The number is at the VM level. If you suspect one of the virtual disks has high latency, use the counter Peak Virtual Disk Read Latency (ms) and Peak Virtual Disk Write Latency (ms).
- **Relationship**
  From the **Related Clusters and Hosts to selected Datastores** widget, select either an ESXi host, a vSphere cluster, or a vSAN cluster. The relevant contention and utilization counters are displayed.

## ESXi Contention Dashboard

The **ESXi Contention** dashboard is the primary dashboard for managing ESXi host performance. The VMware administrator or architect can use it to monitor and troubleshoot any performance issue. If you determine that there is a performance issue, use the **ESXI Utilization** dashboard to see if the cause for the contention is high utilization.

### Design Consideration

The **ESXi Contention** dashboard complements the Cluster Contention Dashboard, and shares the same design consideration.

This dashboard is used as part of your Standard Operating Procedure (SOP). It is designed to be used daily, hence the views are set to show data in the last 24 hours. The dashboard provides performance metrics for virtual machines in the selected data center.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

**How to Use the Dashboard**

- ESXi CPU Performance and ESXI Memory Performance.
  - Review the two distribution charts for an overview of all the ESXi host's utilization and memory performance.
  - Both charts are using the percentage of VM facing performance counter and not the worst performance among VM counter because you are looking at the ESXi performance and not at the single VM performance. See how it handles all the VMs.
  - The bar chart is color coded. Keep the percentage of the VM population not being served under 10%.
- ESXi Hosts Performance.
  - The ESXi hosts performance table lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table is displays the green color, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance greater than 24 hours are likely to be irrelevant.
  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select an ESXi host from the table.
  - All the health charts show the KPI of the selected cluster.
  - For performance, it is important to show both depth and breadth of a performance problem. A problem that impacts one or two VMs require a different troubleshooting than a problem that impacts all VMs in the cluster.
  - Worst CPU overlap among VMs in the host is included as it indicates a lot of interruptions. A running VM might get interrupted because the VMkernel needs the physical core to run something else. High and frequent numbers of interruptions are not healthy and can impact the VM performance.
  - Expect the network error to be 1% and dropped packet to be 0 most of the times, if not always. If it is not zero, analyze it to see if there are any patterns across all ESXi hosts, and bring it up with your network team.

**Points to Note**

- Consider adding a third distribution chart and display the CPU co-stop counter in this third chart, as it complements the CPU ready counter. If your environment has relatively slow network and storage IO, you can add IO wait too.
- Unlike the **Cluster Performance** dashboard, there is no average ESXi hosts performance (%) at the vSphere World level. The reason is most ESXi hosts are part of a cluster and monitoring should be done at the cluster level.
- Certain settings such as power management and hyper threading can impact the performance. Consider adding a property widget to show relevant properties of a selected ESXi host.

### ESXi Utilization Dashboard

The VMware administrator uses the **ESXi Utilization** dashboard with the **ESXi Contention** dashboard to manage performance.

**Design Considerations**

The **ESXi Utilization** dashboard supports the **ESXi Contention** dashboard. Use it to identify vSphere clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, VMware Cloud Foundation OperationsVMware Aria Operations has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

The dashboard complements the Cluster Utilization dashboard, by providing the extra details. Hence it has a similar layout.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

**How to Use the Dashboard**

- ESXi Hosts Utilization.
  - It lists all the ESXi hosts, sorted by the highest utilization in the last one week. If the table is all displaying the green color, then there is no need to analyze further.

  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select an ESXi host from the table.
  - All the utilization charts display the key utilization metrics of the selected cluster.
  - For memory, the high utilization counters are explicitly shown, for example balloon, compressed, or swapped. You might notice they exist even though utilization is not even at 90%, indicating that there was a high pressure in the past. If you look at only utilization, you might think you are safe.
  - The disk IOPS and the disk throughput are split into read and write to gain an insight into the behavior. Some workload is read oriented, while others are write oriented.
  - The network throughput is split into sent (transmit) and received to gain insight into the behavior. The total usage can be misleading because it sums up the send and receive traffic. In reality the network pipe is one for each direction (due to the full duplex nature of Ethernet), and not shared.

**Points to Note**

If your operations team have some forms of standardization that the utilization should not exceed a certain threshold, you can add the threshold into the line chart. The threshold line helps less technical teams as they can see how the real value compares with the threshold. For more information, see the points to note in the ESXi Contention Dashboard.

### Network Performance Dashboard

Use the **Network Performance** dashboard to view performance problems related to network such as high latency, frequent retransmit, and many dropped packets. This dashboard is designed for both the VMware administrator and the Network administrator, to foster a closer collaboration between the two teams.

**Design Considerations**

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

The dashboard activates you to drill down from the distributed switch to the ESXi host and port groups in the switch, and then to the VM.

**How to Use the Dashboard**

- Distributed Switches.
  - The distributed switches table lists all the switches, sorted by the highest packet dropped. The table splits the incoming traffic and the outgoing traffic for better analysis.
  - As the focus is on performance and not capacity, the throughput counters are not shown.
- Select a switch from the distributed switches table.
  - The health chart shows the dropped packet trend over time.
  - It does not narrow down the list of port groups automatically, as the list of port groups are always showing all the port groups in your environment.
  - If necessary, expand the two collapsed widgets. They show the network throughput and broadcast packets. Utilization is also shown so that you can correlate and understand whether the dropped packets are due to higher utilization.
- Port Groups and ESXi Hosts in the selected switch.
  - They get listed when you select a switch from the distributed switches table.
  - Just like the distributed switch, you can also see their relevant counts.
- If your environment has unused network switches, you can filter them out from this list, as this dashboard focuses only on performance.

**Points to Note**

- vSphere network is by nature distributed. Each ESXi contributes to the physical NIC. This represents the physical capacity. Distributed switch and its port groups span across these independent network cards. This makes it harder to define and measure its performance. An unbalance can happen among ESXi hosts or physical NIC. In a sense, it

is like distributed storage ( vSAN). Capacity management does not apply to a port group, since its upper limit (also known as the physical capacity) can vary by even a minute.
- Latency within a data center should be below 1 millisecond. Use VMware Aria Network Insight to study the latency or the retransmitting problems, caused by moving into the lateral traffic.
- Add a physical network using the appropriate management pack.

Most packets are unicast, between a pair of sender and receiver. If your environment has many VMs sending broadcast packets to everyone and multicast packets to many targets, add a Top-N widget to find out which VMs are sending these packets.

### vSAN Contention Dashboard

The **vSAN Contention** dashboard is the primary dashboard for managing vSAN performance. The VMware administrator or architect can use it to monitor and troubleshoot the vSAN cluster performance. If you determine that there is a performance issue, use the **vSAN Utilization** dashboard to see if the cause for the contention is high utilization.

#### Design Considerations

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

The **vSAN Contention** dashboard complements the vSphere Cluster Capacity, and shares the same design consideration. It focuses on the storage and vSAN specific metrics, and does not repeat what is already covered. It does not list any non vSAN cluster.

#### How to Use the Dashboard

- vSAN Peak VM Latency, vSAN Peak CPU Ready, vSAN Peak Dropped Packet.
  - Review the three distribution charts for an overview of all the vSAN clusters performance.
  - The vSAN peak VM latency chart shows the distribution of disk latency experienced by all the VMs in the cluster. You should expect most of the VMs to experience latency that matches your expectation. For example, in an all flash systems, the VMs should not have >20 ms disk latency. If your vSAN environment is all flash, you must adjust the distribution bucket to a more stringent set.
  - The vSAN peak CPU ready chart shows if any of the vSAN kernel modules has to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not wait for CPU time. vSAN gets higher priority than VM World as it lives in the kernel space.
  - The vSAN peak dropped packet chart shows if any of the vSAN clusters are dropping packet in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.
- vSAN Clusters.
  - It lists all the vSAN clusters, sorted by the least performing.
  - It lists all the ESXi hosts, sorted by the worst performance in the last 24 hours. If the table is showing all green, then there is no need to analyze further. The reason 24 hours is selected instead of one week is that the performance issues greater than 24 hours are likely to be irrelevant.
  - You can change the time period to the period of your interest. The maximum number is reflected accordingly.
- Select a vSAN cluster from the vSAN clusters table.
  - All the health charts show the KPI of the selected cluster.
  - If you are using SMART, the two heat maps at the bottom of the dashboard provide early warning.

#### Points to Note

- A large vSAN cluster can have many components. Each of these components can have multiple performance metrics. The total number of KPI can reach hundreds of metrics. For example, take a 10 node cluster. It can have 530 counters to check. VMware Cloud Foundation OperationsVMware Aria Operations aggregates them by introducing a set of KPIs. This analysis reduces the number to a more manageable number. The following table shows the KPIs and their formula.

| Name | What it is |
|---|---|
| Max Capacity Disk Latency (ms) | Highest latency among all capacity disks take the worst, not average, as the latency in a single capacity disk is already an average of all its VMs. If there are 50 VMs on the disk and 30 are issuing IO on it, then its average is among 30. |
| Min Disk Group Write Buffer Free (%) | Lowest free capacity among all the disk group write buffers. If this number is low, one of your buffers is not enough. While you want to maximize your cache, a low number is an early warning for capacity management. |
| Max Disk Group Read Cache/Write Buffer Latency (ms) | Each disk has a Read Cache Read Latency, Read Cache Write Latency (for writing into cache), Write Buffer Write Latency, and Write Buffer Read Latency (for de-staging purpose). This takes the highest among all these four numbers and the highest among all disk groups. It is the max of the max because each of the four datapoints is an average of all the VMs on it. |
| Sum Disk Group Errors | Sum of the bus reset + sum of commands canceled among all the disk groups. You must use sum and not get the max as each member should return zero. |
| Count Disk Group Congestion Above 60 | The number of disk groups congestion greater than 60. 60 is hardcoded in the vSAN Management Pack as it is a good starting point. As any congestion above 60 serves an early warning, count how many of such occurrences happen. |
| Max Disk Group Congestion | The highest congestion among all disk groups. A high number indicates that at least one disk group is not performing. |
| Min Disk Group Capacity Free (%) | The lowest free capacity among all disk groups. A low space triggers rebalance. |
| Min Disk Group Read Cache Hit Rate (%) | The lowest hit rate among the disk group read cache. Ensure that this number is high as it indicates that the read is served by cache. |
| Sum vSAN PortGroup Packets Dropped (%) | Sum of all vSAN VMkernel port RX dropped packet + TX dropped packet. You should expect no dropped packet in your vSAN network. |

### vSAN File Services

The VMware administrator uses the **vSAN File Services** dashboard to monitor the file services running in their vSAN environment.

#### Design Considerations

This dashboard is designed to complement the vSAN file services management provided by the vCenter. The vCenter is more of an administrative tool, while VMware Cloud Foundation OperationsVMware Aria Operations is more of an operations tool. Each tool performs their specific functions and does not duplicate information.

#### How to Use the Dashboard

- File Shares by Used Space and Latency.
  - Review the file shares by used space and latency heat map.
  - It shows all the file shares in your environment.

  - The greater the use (consumption), the greater the box, so you can easily see the most consumed ones.
  - The file shares are colored by latency. You must watch out for boxes with red color.
- vSAN Clusters with File Services activated.
  - It lists all the vSAN clusters with file services activated, giving a convenient view to see which clusters have these settings turned on.
- Select a vSAN cluster from the vSAN clusters with file services activated table.
  - The file servers in the selected vSAN cluster are shown. When you select a file server, it filters the file shares list to show the file shares in the selected file server.
  - The file shares in the selected vSAN cluster are shown. Selecting a file share displays all the relevant KPI on the file share.

**Points to Note**

vSAN File Servers and vSAN File Shares are two new objects in VMware Aria Operations Management Pack for vSAN.

**vSAN Performance Dashboard**

Use the **vSAN Performance** dashboard along with the **Cluster Capacity** dashboard.

**How to Use the Dashboard**

The **vSAN Performance** dashboard is organized into sections for ease of use.

- **Cluster Analysis**
  The **vSAN Clusters** widget lists all the vSAN clusters, starting with cluster with the highest VM disk latency. By default, the widget displays data from the last 24 hours.

  The first column displays if the distribution of disk latency is experienced by all the VMs in the cluster. You can expect a majority of the VMs to experience latency that matches your expectation.

  The second column displays if any of the vSAN kernel modules have to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not be waiting for CPU time.

  The third column displays if any of the vSAN clusters are dropping packets in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.

  Select a cluster to investigate further. The VM latency distribution is automatically displayed.
- **Contention**
  You can view various disk-related contention counters of the cluster from the **Contention** widget.
- **Utilization**
  Contention metrics are complemented by utilization metrics. A large block size can result in high throughput in a relatively low IOPS. If you see a large block size when you are not expecting it, investigate which applications are using it.
- **Disk Groups**
  You can drill down to the disk group level. All the counters are the worst value among the disk groups.
- **Read Cache**
  All the values displayed are the worst values among the read cache of the disk group.
- **Other KPIs**
  A problem in performance can also be caused by non-storage. vSAN resync is a type of utilization metric, but its presence can impact performance.
- **Disk Group Analysis**
  You can drill down to individual disk groups from the **Disk Groups** widget. Ensure that the disk groups are fairly balanced. Select the disk group that you want to analyze. Both the contention and utilization metrics are automatically displayed.
- **Cache Disks**

You can drill down to cache disks from the **Cache Disks** widget. Ensure that the configuration is consistent. Select the cache disk you want to analyze. Both the contention and utilization metrics are automatically displayed.

- **Capacity Disks**
  You can drill down to capacity disks from the **Capacity Disks** widget. Ensure that the configuration is consistent. Select the capacity disk you want to analyze. Both the contention and utilization metrics are automatically displayed.

### vSAN Utilization Dashboard

The VMware administrator uses the **vSAN Utilization** dashboard with the **vSAN Contention** dashboard to manage performance.

#### Design Consideration

The **vSAN Utilization** dashboard supports the **vSAN Contention** dashboard. Use it to identify vSAN clusters with high utilization in a selected data center. When utilization exceeds 100%, performance can be negatively impacted, especially when a VM experiences contention. By default, VMware Cloud Foundation OperationsVMware Aria Operations has a 5-minute collection interval. For 5 minutes, there might be 300 seconds worth of data points. If a spike is experienced for a few seconds, it might not be visible if the remaining 300 seconds is low utilization.

To view the common design considerations among all performance management dashboards, see the Performance Dashboards.

#### How to Use the Dashboard

- Clusters Utilization.
  - It lists all the vSAN clusters, sorted by the least performing.
- Select a vSAN cluster from the clusters utilization table.
  - All the health charts show the KPI of selected cluster.
- Disk Groups
  - It lists all the vSAN clusters, sorted by the least performing.
- Select a Disk Group from the disk groups table.
  - All the health charts show the KPI of selected cluster.

#### Points to Note

- The **vSAN Utilization** dashboard complements the **vSAN Contention**. For more information, see the points to note in the vSAN Contention Dashboard.

### vSAN ESA Performance Dashboard

The vSAN ESA (Express Storage Architecture) Performance dashboard provides an overview of the performance issues related to your vSAN clusters. The dashboard displays the read and write latency issues that affect the vSAN clusters present in your environment. You can use this dashboard to identify and troubleshoot performance issues related to Contention and Utilization across VMs that are part of the vSAN cluster.

#### How to Use the Dashboard

The **vSAN ESA Performance** dashboard is organized into sections for ease of use.

- **Cluster Analysis**
  The **vSAN ESA Clusters** widget lists all the vSAN clusters, starting with cluster with the highest VM disk latency. By default, the widget displays data from the last 24 hours.

  The first column displays if the distribution of disk latency is experienced by all the VMs in the cluster. You can expect a majority of the VMs to experience latency that matches your expectation.

  The second column displays if any of the vSAN kernel modules have to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not be waiting for CPU time.

The third column displays if any of the vSAN clusters are dropping packets in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.

Select a cluster to investigate further. The VM latency distribution is automatically displayed.

- **Contention**
  You can view various disk-related contention counters of the cluster from the **Contention** widget.
- **Utilization**
  Contention metrics are complemented by utilization metrics. A large block size can result in high throughput in a relatively low IOPS. If you see a large block size when you are not expecting it, investigate which applications are using it.
- **Configuration**
  You can view the configuration details such as storage type, number of VMs, ESXi, Storage Pools, Physical Disks of the selected vSAN cluster.
- **Storage Pools**
  You can view the configuration, contention, and utilization details of the Storage Pool for the selected vSAN cluster.
- **Physical Disks**
  You can view the configuration, contention, and utilization details of the Physical Disks for the selected vSAN cluster.
- **Related Objects to Selected vSAN Cluster**
  You can view the details of the related objects for the selected vSAN cluster.

### vSAN OSA Performance Dashboard

The vSAN OSA (Original Storage Architecture) Performance dashboard provides an overview of the performance issues related to your vSAN clusters. The dashboard displays the read and write latency issues that affect the vSAN clusters present in your environment. You can use this dashboard to identify and troubleshoot performance issues related to Contention and Utilization across VMs that are part of the vSAN cluster.

#### How to Use the Dashboard

The **vSAN Performance** dashboard is organized into sections for ease of use.

- **Cluster Analysis**
  The **vSAN Clusters** widget lists all the vSAN clusters, starting with cluster with the highest VM disk latency. By default, the widget displays data from the last 24 hours.

  The first column displays if the distribution of disk latency is experienced by all the VMs in the cluster. You can expect a majority of the VMs to experience latency that matches your expectation.

  The second column displays if any of the vSAN kernel modules have to wait for CPU. Expect this number to be near 0% and below 1%, as vSAN should not be waiting for CPU time.

  The third column displays if any of the vSAN clusters are dropping packets in the vSAN network (not the VM network). vSAN relies on the network to keep the cluster in-sync. This number should be near 0% and less than 1%.

  Select a cluster to investigate further. The VM latency distribution is automatically displayed.
- **Contention**
  You can view various disk-related contention counters of the cluster from the **Contention** widget.
- **Disk Group Analysis**
  You can drill down to individual disk groups from the **Disk Groups** widget. Ensure that the disk groups are fairly balanced. Select the disk group that you want to analyze. Both the contention and utilization metrics are automatically displayed.
- **Utilization**
  Contention metrics are complemented by utilization metrics. A large block size can result in high throughput in a relatively low IOPS. If you see a large block size when you are not expecting it, investigate which applications are using it.

- **Configuration**
  You can view the configuration details such as storage type, number of VMs, ESXi, Storage Pools, Physical Disks of the selected vSAN cluster.
- **Compute and Network**
  You can view the compute and network metrics for the selected vSAN cluster.
- **Cache Disks**
  You can drill down to cache disks from the **Cache Disks** widget. Ensure that the configuration is consistent. Select the cache disk you want to analyze. Both the contention and utilization metrics are automatically displayed.
- **Capacity Disks**
  You can drill down to capacity disks from the **Capacity Disks** widget. Ensure that the configuration is consistent. Select the capacity disk you want to analyze. Both the contention and utilization metrics are automatically displayed.
- **Related Objects to Selected vSAN Cluster**
  You can view the details of the related objects for the selected vSAN cluster.

#### vSphere Performance Profiling Dashboard

Use the vSphere Performance Profiling dashboard to determine a suitable Service Level Agreement (SLA).

This dashboard displays the performance of all vSphere clusters and their shared datastores. For each VM, the dashboard displays the performance metric at the 99th percentile. For each SLA (CPU, Memory, Disk), the dashboard shows the worst and the average performing VM.

> **NOTE**
> Avoid setting the SLAs at a 100th percentile, as the 100th percentile usually indicates an outlier.

## Private AI (GPU) Dashboards

Use the Private AI (GPU) dashboards to monitor and troubleshoot GPU issues in VMware Aria OperationsVMware Cloud Foundation Operations.

To access the dashboards, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, navigate to **All › Private AI**.

The following dashboards are available:

| Dashboard Name | Purpose |
|---|---|
| GPU Equipped Clusters | Use this dashboard for the details related to GPU compute utilization and GPU memory usage at the cluster level, host level, and GPU level. |
| GPU Overview | Use this dashboard to view if any GPUs have high-temperature. This dashboard also highlights GPUs with low to zero usage by analyzing their capacity based on compute and memory utilization. |

#### GPU Equipped Clusters Dashboard

Use the GPU Equipped Clusters dashboard to view the the clusters that have GPU activated devices. You can view the GPU device information at the cluster, host, and GPU level.

This dashboard displays the total number of GPU devices, GPU compute usage, memory usage, and memory at the cluster and host level. You can select a GPU to view the cluster level GPU utilization and the host level GPU utilization. You can also view the GPU metrics and properties at the GPU level, and further drill down to get an overview of the GPU temperature and GPU instance metric chart.

### GPU Overview Dashboard

Use the GPU Overview dashboard to view the GPU summary of all the GPUs present in your environment.

This dashboard displays the total number of GPUs and highlights GPUs as per heat map, memory usage (%) and compute utilization (%). You can select a GPU to view its relationship topology, the GPU instance name and properties, and the GPU instance metric chart at the GPU level.

## Sustainability Dashboards

Using the Sustainability dashboards in VMware Aria Operations, you can view the reduction in carbon footprint achieved using virtualization, you can view clusters and data centers and compare them based on CO2 emissions or power consumption, you can identify idle VMs, and also identify old hardware, such as compute and storage.

The following Sustainability dashboards are added to the predefined VMware Aria Operations dashboards:

- Carbon Efficiency with Virtualization
- Carbon Transparency
- Environmental Impact of Idle VMs
- Green Supply

### Carbon Efficiency with Virtualization Dashboard

Use this dashboard to view the reduction in carbon footprint achieved using virtualization and also the reduced number of servers, reduced power consumption, and reduced carbon emission achieved using virtualization.

**How to Use the Dashboard**

- The banner at the top of the dashboard is a sample image, and you can upload your company logo, tag line, and so on.
- The **Greener Planet Contribution** widget provides details of power saved for the vSphere World object in percent and compares the power consumption of a vCenter environment before and after virtualization assuming that 100W is the power consumption for a low-range server before virtualization.
- Use the **CO2 Emissions Saved** and **Electricity Cost Saving** widgets to view the sum of carbon emissions and electricity saved for the last 24 hours.
- Using the **Power Consumption** widget, you can view power consumed before and after virtualization. Power consumption before virtualization is based on the assumption that a low-range server consumes 100W. Hence, power consumption before virtualization is calculated using the count of VMs, which is the number of physical servers if not virtualized, at 100W per server.
- Using the **CO2 Emission** widget, you can view CO2 emissions before and after virtualization. CO2 emission per kWh is assumed as 0.709 kg (Based on reference values from Greenhouse Gas Equivalencies Calculator). The value is calculated in Kilogram (kg).
- The **CO2 Emission Chart** and the **Power Consumption Chart** provides CO2 emissions (in kg) for the last 24 hours and power consumption (in kWh) for the last 24 hours respectively.

**Custom Values**

- **Predefined Values**
  - Power consumption of a small server (1 socket, 10 cores, 32 GB RAM) = 0.1 kW (Assumption used in calculating power consumption before virtualization).
  - CO2 emission per kWh = 0.709 kg (Reference values from Greenhouse Gas Equivalencies Calculator).
  - Electricity Cost per kWh = $0.108 (Reference from VMware TCO Reference Calculator).
- You can add a custom property for CO2 emission and electricity cost which can be applied per cluster compute resource.

Custom values of CO2 emission per kWh can be added to each cluster compute resource by creating a custom property with the name **CO2 Emission** and type as **Numeric** and the relevant custom values. You can add custom values for Cost of Power to each cluster compute resource by creating a custom property with name **Electricity Rate** and type as **Numeric** and the relevant custom values.

> **NOTE**
> If you do not define custom properties, out of the box values for CO2 emission and electricity costs are used.

**Points to Note**

Your actual savings may be much more. The following are not included:

- Physical buildings and land. With virtualization, you consume less carbon foot print. This results in less physical rack.
- Network equipment. Fewer number of physical servers results in lesser network ports. Because firewall, load balancers, IDS, and IPS can be virtual machines, you have less equipment.
- Other components like UPS, lighting, cooling, and labor.

**Carbon Transparency Dashboard**

Use this dashboard to view clusters and data centers and compare them based on CO2 emissions or power consumption. You can then identify the most green cluster to provision workloads. You can compare the power consumption of each compute component in the data center, showcase all the compute components with the lowest power consumption, compare physical data centers based on power consumption, and compare hardware models to find optimal power consumption.

**How to Use the Dashboard**

As larger clusters consume more power than smaller clusters, the total power consumption cannot be used to determine a cluster to provision the next workload. Hence Power Efficiency is calculated based on the power consumption per GHz of CPU usage.

- The banner at the top of the dashboard is a sample image, and you can upload your company logo, tag line, and so on.
- Use the **Top-10 Green Clusters by Power Efficiency** table to view clusters that are power efficient. Power efficiency is calculated as power usage per GHz of CPU Usage. Power efficiency alone cannot be used as the criteria to provision a VM, capacity must also be available in the cluster. View the `Capacity Remaining%` and `Time Remaining` information for each cluster to make an informed decision.
- Use the **Top-10 Green Clusters by Power Consumption** table to view clusters that consume less power. Use this table when you compare clusters based on the total power consumption. Larger clusters will consume more power than smaller clusters, but may still not be the most efficient cluster based on power efficiency. Power consumption alone cannot be used as the criteria to provision a VM, capacity must also be available in the cluster. View the `Capacity Remaining%` and `Time Remaining` information for each cluster to make an informed decision.
- The goal of the **Carbon Transparency** dashboard is to help you identify the greenest cluster to provision the next VM. However, the most green cluster may not have enough capacity to provision a VM, and hence you also need to check capacity before you decide on the greenest cluster to provision. Use relevant metrics for the selected cluster which can help you identify a target cluster. The metrics are: Power Consumption of the Cluster (kWh), CPU Usage% in the cluster, and Memory usage% in the cluster. Use the associated widgets to view the metrics. The metrics are available for the last 30 days to help you identify the target green cluster to provision the next VM.

- Use the **Heatmap of vSphere Clusters based on Power Consumption** to view all the clusters based on Power Consumption in kWh. The size of each cluster in the heatmap is determined by the number of VMs in the cluster.
- When you add geo tags to clusters or physical data centers, those objects can be mapped in the **Geo** widget. Ensure that you add geo tags to the relevant objects.

### Environmental Impact of Idle VMs Dashboard

Use this dashboard to identify idle VMs in the data center, power off/delete idle VMs to reduce power consumption and improve data center efficiency, and identify remediation plans to offset the damage caused by those idle VMs which cannot be removed.

#### How to Use the Dashboard

- The banner at the top of the dashboard is a sample image, and you can upload your company logo, tag line, and so on.
- The **Cluster's Idle VM** table provides a list of clusters with reclaimable memory ,vCPUs, and disk space from idle VMs. You can also view the total resources that can be reclaimed from idle VMs. Click a cluster row to view the relevant sustainability metrics.
- Use the **Relevant Metrics of Selected Cluster** widget to view metrics of the selected cluster such as, power wasted by idle VMs (Wh), CO2 emissions by idle VMs (kg), trees to offset CO2 emission of idle VMs.
- The **Wasted Power from Idle VMs** heatmap allows you to view vSphere clusters configured to display wasted power from idle VMs. Power in watt-hour (Wh) is used to build the heatmap. A threshold value of 800 Wh changes the heatmap to red.
- The **CO2 Emission from Idle VMs** heatmap allows you to view vSphere clusters configured to display CO2 Emission from idle VMs. CO2 emission in kilogram (kg) is used to build the heatmap. A threshold value of 750 kg changes the heatmap to red.
- The **Trees Required to Compensate** heatmap allows you to view vSphere clusters configured to display the number of trees to be planted to compensate for the damage caused by idle VMs. A threshold value of 70 trees changes the heatmap to red

#### Custom Values

- **Predefined Values**
  - Tree offset for CO2 Emission = 16.511 kg (36.4 pounds of carbon per tree. Refer to Greenhouse Gases Equivalencies Calculator) which is equivalent to 36.4/2.2046 kg of carbon per tree).
  - CO2 emission per kWh = 0.709 kg ( Reference values from Greenhouse Gas Equivalencies Calculator).
- Custom values for CO2 emission can be applied per cluster compute resource by adding a custom property. Custom values of CO2 emission per kWh can be added to each cluster compute resource by creating a custom property with the name **CO2 Emission** and type as **Numeric** and the relevant custom values. Custom values of tree offset to compensate CO2 emission can be added to each cluster compute resource by creating a custom property with the name **Trees to Offer** and type as **Numeric** and the relevant custom values.

  **NOTE**
  If no custom properties are defined, out of the box values for CO2 emission are used.

### Green Supply Dashboard

Use this dashboard to identify old hardware, such as compute and storage, in the data center and replace them with new generation hardware components that are power efficient. You can also use this dashboard to reduce the overheads and buffers and identify smaller clusters that have higher overheads. The aim is to run with fewer overheads and buffer, without compromising on performance.

**How to Use the Dashboard**

- Smaller clusters have a relatively higher overhead. A cluster with two nodes has 50% overhead, while a cluster with 10 nodes has only 10% overhead. Clusters with lesser capacity require more hosts and hence consume more electricity. The **Small Clusters** table lists clusters that meet one of the following criteria:
  - <=4 nodes
  - <=120 CPU cores and < 1 TB memory

  Click a cluster row to view the context of the selected cluster. An empty widget indicates that the defined green goals are met.
- Advancements in technology help ESXi hosts to deliver higher efficiency. ESXi hosts can deliver more CPU and memory capacity, often with low power requirements. The **Ageing Compute Hardware** table lists ESXi hosts that meet one of the following criteria:
  - ESXi version 6.0 or older
  - <=40 CPU cores and < 256 GB of memory

  Click a cluster row to view the context of the selected cluster. An empty widget indicates that the defined green goals are met.
- Just like compute hardware, newer storage hardware is more power efficient than older storage hardware. The **Ageing Storage Hardware** table lists datastores that meet the following criteria:
  - VMFS version 5 or older.
  - Not a local datastore.

  Click a cluster row to view the context of the selected cluster. An empty widget indicates that the defined green goals are met.

  If you have goals that are different from those defined, you can modify the criteria of the widgets by updating the filters.

## Service Discovery Dashboards

Using the service discovery dashboards, you can determine the inter-dependencies of virtual machines and the dependencies of each service in the respective virtual machines.

The following service discovery dashboards are added to the predefined VMware Aria OperationsVMware Cloud Foundation Operations dashboards:

- Service Distribution
- Service Relationships
- Service Visibility
- Virtual Machine Relationships

### Service Distribution Dashboard

You can use the dashboard to view the distribution of different services in the selected data center, cluster, or a host system. You can also view known and unknown services including the category and distribution percentage across a vSphere resource.

You can use the dashboard widgets in several ways:

- **Inventory Item**: Use this widget to view a hierarchical representation of objects in the form of badges.
- **Known Services Distribution**: Use this widget to view different services discovered from a selected object.
- **Service Categories**: Use this widget to view the service categories that are discovered by selecting an object from the resource widget.
- **User Defined Services Distribution**: Use this widget to view a list of user-defined services.

### Service Relationships Dashboard

You can use the dashboard to view properties of the service such as the install path, the ports used, and the version. You can also view the relationship between the services that run on other VMs.

You can use the dashboard widgets in several ways:

- **List of Services Discovered**: Use this widget to view the services that have been discovered.
- **Connections from the Selected Services**: Use this widget to view the relationship between the services and the other services running on the VMs.
- **Properties of the Selected Service**: Use this widget to view the properties of the selected services.

#### Service Visibility Dashboard

You can use the dashboard to view a list of VMs without service visibility and VMs with user-defined services after you select a vSphere object.

You can use the dashboard widgets in several ways:
- **Inventory Tree**: Use this widget to view a hierarchical representation of objects in the form of badges.
- **Virtual Machines without Service Visibility**: Use this widget to view information about services where discovery has failed.
- **Virtual Machines with User-Defined Services**: Use this widget to view a list of VMs where the user has defined such services.

#### Virtual Machine Relationships Dashboard

You can use the dashboard to view a list of VMs with service discovery details such as, status, method, incoming/outgoing connections, and protection groups. When you select a VM, the dashboard displays a list of discovered services on the VM, the relationships of the VMs with other VMs based on the relationships of the discovered service.

You can use the dashboard widgets in several ways:
- **List of virtual machines**: Use this widget to view all the VMs discovered by the vCenter.
- **Node relationship of the selected VM**: Use this widget to view the relationship between the objects.
- **List of Services running in the selected VM**: Use this widget to view all the properties of the selected VM.
- **Connections of Virtual Machines**: Use this widget to view the relationship between one or more VMs.

## Skyline Operational Overview Dashboard

The Skyline Operational Overview dashboard displays your inventory summary, high level configuration data organizing by attributes such as vCenter instances, ESXi hosts, virtual machines, and VM Guest Tools configurations.

#### Cluster Configuration

Use this dashboard to view the cluster configurations and settings related to VMware Distributed Resource Scheduler (DRS), VMware High Availability (HA), and HA Admission Control Enabled. For example, to view how many hosts have HA or DRS configured, you can click the pie chart to view the cluster information. Click a cluster to view the cluster summary, alerts, and other information.

#### ESXi Host Configuration

You can view your ESXi host configuration organized by hardware model, BIOS version, and host version. Click the pie chart to view the host information. Click any host to view the cluster summary, alerts, and other information.

#### Virtual Machine Configuration

You can view your virtual machine capacity in terms of CPU, RAM, and disk size. Clicking the pie chart displays the VM information. Click any VM to view the cluster summary, alerts, and other information.

#### VM Guest OS Distribution

Monitor your guest tool information such as OS distribution, tool version, and tool status. Clicking the pie chart displays the tools information. Click any tool to view the summary, alerts, and other information.

# Dashboard Library

### Executive Summary Dashboards

The requirements of the CIO, Head of Global Infrastructure, and IT Senior Management vary from the requirements of the technical teams. The **Executive Summary** dashboards provide an overall information on capacity and inventory in business terms.

These dashboards allow you to display problems related to budget and resource, and provide visibility to the senior management into the live environment. By doing this, you can prove the need for additional hardware. If there is wastage that has to be reclaimed, you can display where and how large the wastage is using these dashboards. VMware Aria OperationsVMware Cloud Foundation Operations provides two example dashboards to get you started. As each executive might have a unique requirement or preference, the dashboards can be customized accordingly. The five principles displayed in the following figure are used to design the **Executive Summary** dashboards.

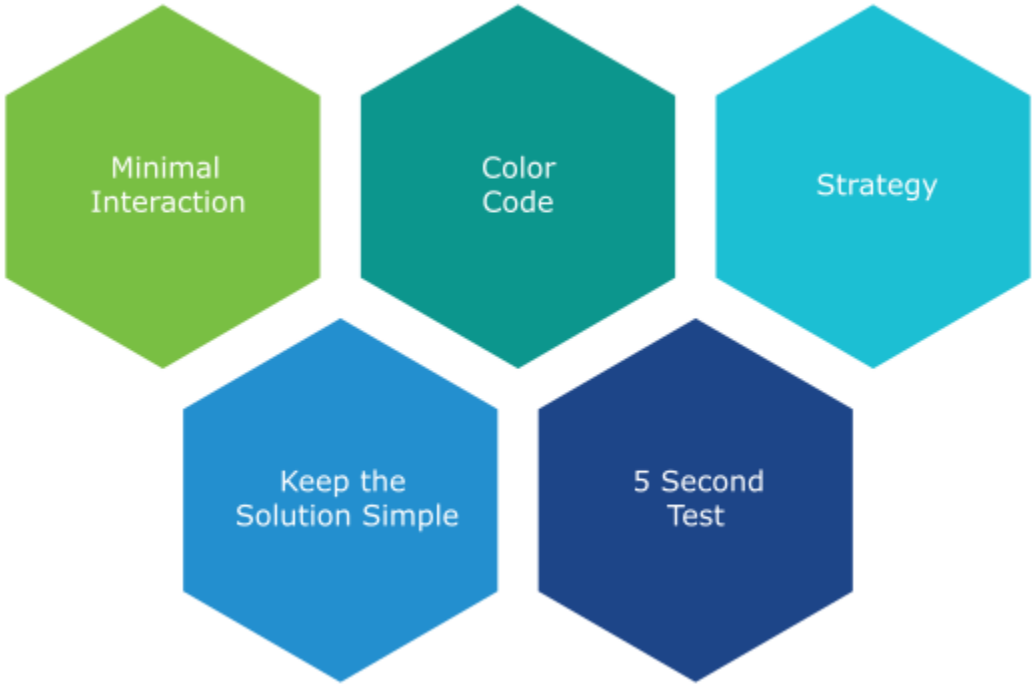


- Keeping the interaction, such as clicking, zooming, and sorting to a minimal.
- Use of color codes to have a user interface that is easy to understand.
- Each dashboard answers a specific question and the information is presented in business terms.
- Keep the solution simple and have a portal that is easy to access.
- Ensure that the dashboards are understood within five seconds.

### Capacity Summary Dashboard

The **Capacity Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Inventory Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

#### Design Considerations

See Executive Summary Dashboards for common design considerations among all the dashboards for the IT senior management.

### How to Use the Dashboard

The **Capacity Summary** dashboard has two sections:

- The top section of the dashboard provides a summary at the vSphere World level.
  - The **VM Growth** widget displays the weekly average of the VM growth and provides holistic visibility of overall growth across all data centers for both running and powered off workloads. If an increase in the VM count is not accompanied by a corresponding increase in utilization, these newly provisioned VMs are likely not yet used.
  - The **Overcommit Ratio** widget highlights the efficiency gained by vSphere virtualization running multiple workloads on a shared infrastructure. Overcommitment has to be further reviewed along with elevated resource contention to understand the impact of performance on VMs competing for resources. In general, Overcommit is required to be financially more economical than the public cloud.

    **NOTE**
    VMware Aria OperationsVMware Cloud Foundation Operations uses physical CPU Cores not Logical Cores (Hyper-threading) for all CPU-based capacity calculations.

- The bottom section of the dashboard activates drill down into individual compute or storage capacity.
  - Capacity is split into Compute (vSphere Clusters) and Storage (Data stores) views. The heat map displays capacity by size and color by time remaining. By selecting either Clusters or Data stores, you can further drill down understand the remaining capacity and time (in days).

### Points to Note

- Capacity remaining is not displayed at the vSphere World level as it can be misleading, especially in a global, or large infrastructure. Clusters also tend to serve a different purpose and they are not interchangeable.
- If you are using both on-prem and external cloud, for example, VMware on AWS, consider splitting the dashboard into two columns.

## Inventory Summary Dashboard

The **Inventory Summary** dashboard is used by the Ops team to explain capacity to IT Management. This dashboard works together with the **Capacity Summary** dashboard. The inventory provides details on available resources and what is running on these resources. The capacity provides details on the remaining capacity and time.

### Design Considerations

See Executive Summary Dashboards for common design considerations among all the dashboards for the IT senior management.

### How to Use the Dashboard

- The **Summary** widget provides a quick view of the key inventory number.
  - The scoreboard is interactive. This widget drives the eight pie charts that are placed at the bottom of the dashboard. Since all the information is at the vSphere World level, clicking any of them will display details of the total inventory.
- Select any data center from the **Datacenters** widget.
  - This widget drives Clusters and Datastores so that you can quickly view what you have in a given data center and related capacity.
  - For a small environment, the vSphere World is displayed so you can view all the VMs in the environment.
  - To sort by any of the columns in the table, click the column title.
- The eight charts in the dashboard provide details of the inventory. They are driven by **Datacenters**, **Compute**, **Storage**, and **Summary** widgets.

**Points to Note**

- Understand the relationship hierarchy in vSphere. For example, Compute (cluster) is not a parent of Storage (datastore), so logically it is not possible to display datastores in a cluster. Data center consists of compute (cluster), network (distributed switch), and storage (datastore).
- Datastores do not drive the pie chart. This is a known limitation in the View widget.
- If your senior management wants to view the largest VM in a given environment, add a Top-N widget to list the top 10 largest consumers so that CPU, memory, and the disk details are highlighted.

**Network Operation Center**

A dashboard projected on the large screen serves a different business purpose than a dashboard on your laptop or desktop. It is placed strategically because it displays a time sensitive information. Dashboards complement alerts and cannot replace it. The five principles displayed in the following figure are used to design the predefined **Network Operation Center** dashboards.

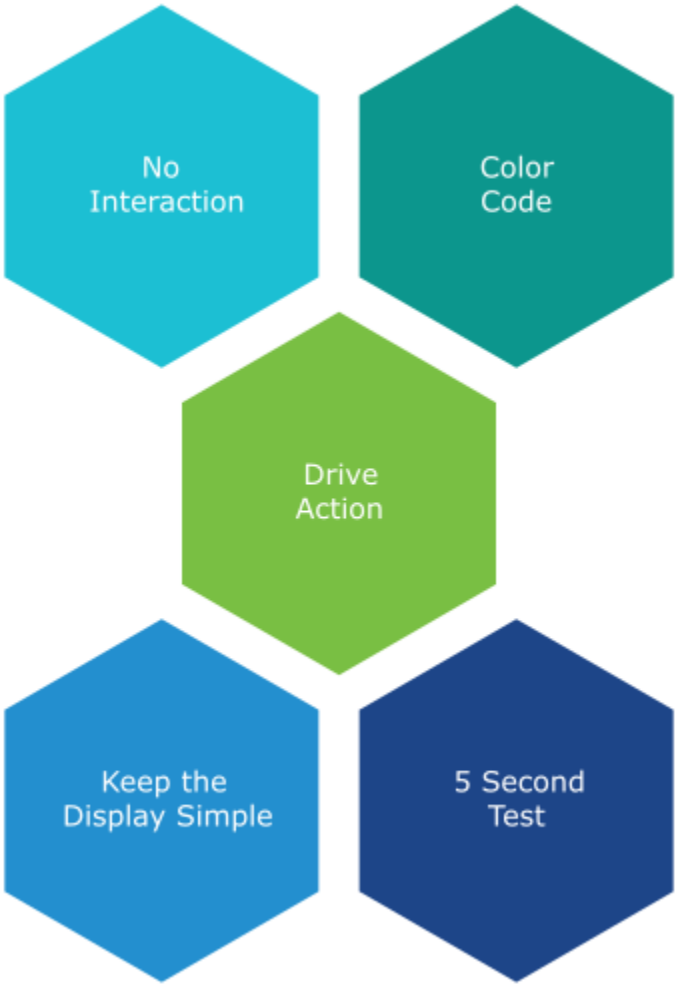


- Keeping the interaction, such as clicking, zooming, and sorting to a minimal. Avoid having buttons, use of mouse or keyboard to view data.
- Use of color codes to have a user interface that is easy to understand.
- Displaying content that drives action. Display of live information as the focus is on immediate remediation. Problems that need immediate actions are displayed, for example, stop provisioning of new VM or take action on VMs that abuse the shared infrastructure.

- Display of problems that do not require immediate attention are avoided, for example, increase supply of infrastructure, such as adding hardware.
- Keep the display simple and have a portal that is easy to access.
- Dashboards are designed to display minimal and critical information only.
- Displays of numbers in percentage, with 0% being poor and 100% being perfect. To display utilisation, you can use the following markers:
  - 50% indicates good and balanced utilization. However, the ideal value is 75%
  - 0% indicates wastage
  - 100% indicates high utilization
- Ensure that the dashboards are understood within five seconds.

**Live! Cluster Performance Dashboard**

The **Live! Cluster Performance** dashboard provides live information on whether the requests of the VMs are met by their underlying compute clusters. This dashboard focuses on CPU, Memory, and the performance of the clusters. Use this dashboard to view if there is any problem in meeting the demands of the VMs and if there is any unbalance within a cluster. The **Live! Cluster Performance** dashboard is the primary dashboard and it complements the **Live! Cluster Performance** dashboard which is the secondary dashboard. This secondary dashboard displays if the performance problem is caused by high utilization. The primary dashboard answers the question 'Is our IaaS performing?', while the secondary dashboard answers the question 'Is our IaaS working hard?'.

**Design Consideration**

The **Live! Cluster Performance** dashboard displays three heat maps. The heat maps complement each other and must be used together. The location of each cluster and ESXi hosts within those clusters is identical in all heat maps. The fixed positioning allows you to compare if the problem is caused by memory contention, CPU ready, or CPU co-stop.

The sizes of each cluster and ESXi hosts are constant. Variable sizing creates a distraction and can result in small boxes, making it difficult to read.

The focus of the performance is on the population and not on a single VM. This is not a single VM troubleshooting dashboard but a dashboard focusing on infra problem. As the infra counter is mathematically an aggregation of VM counters, you must have a right roll-up strategy. As the goal is to provide an early warning, do not use the average as a roll-up technique. Use the percentage of the population exceeding a threshold. The threshold is set to be stringent to receive an early warning.

**How to Use the Dashboard**

Review the heat maps, **Memory Contention**, **CPU Ready**, and **CPU Co-Stop** and see if there is any color other than green.

- Green indicates that almost 100% of the VMs have received the CPU and memory that was requested. The threshold is set such that if the 10% of the VM population does not receive the requested resources, then the heat map turns red.
- Red indicates an early warning. Stringent thresholds are used to activate proactive attention and remediation operations. The heat map can turn red because of the high standard that is applied even when there is no complaint from the VM owner yet.
- The light gray indicates that there is no VM running on the host and the metric is not computing.

View if there is any unbalance.

- There are two types of unbalance, cluster unbalance, and resource type unbalance.
- The ESXi hosts are grouped by the cluster, so that the unbalance within a cluster can be easily viewed. Cluster unbalance is a real possibility and it is best monitored and not just assumed.
- If the three heat maps are different, then there is a resource unbalance. For example, if the memory contention is mostly red, but the two CPU heat maps are green, it means you have an unbalance between memory and CPU.
- If a single ESXi host displays different color across the three heat maps, it indicates that there is an unbalance between the CPU and memory resources in the host.

For NOC Operator, drill-down by selecting one of the VMs on the heat map.
- The **Trends of Selected ESXi Host** widget will automatically display the performance counters. To hide any metric, click the name in the legend.

As part of the deployment, configure auto-rotate among the NOC dashboards. If you want to view one dashboard, then you can remove the VMware Aria OperationsVMware Cloud Foundation Operations menu by using the URL sharing feature. This makes the overall user interface presentable and allows you to focus on the dashboard.

**Points to Note**
- You can add Disk Latency if you have the screen real estate. Use the counter 'Percentage of Consumers facing Disk Latency (%)'. It is a part of a datastore object, not a cluster, as a VM in a cluster can have disks across multiple datastores. Organize this storage performance by data center and not by the cluster.

### Live! Cluster Utilization Dashboard

The **Live! Cluster Utilization** dashboard complements the **Cluster Performance** dashboard. Use this dashboard to view the clusters that are working excessively and are close to their physical limit. This dashboard displays ESXi hosts that have CPU or memory saturation that can lead to performance issues for the VMs running on the host.

**Design Considerations**

This dashboard is designed to complement the **Live! Cluster Performance** dashboard and it shares design considerations.

**How to Use the Dashboard**

As this dashboard has an identical design with the **Live! Cluster Performance** dashboard, it has the same usage procedure. Unlike the heat maps in the **Live! Cluster Performance** dashboard, the three heat maps in this dashboard have a different scale, reflecting the different nature of the counters.

Logically, memory is a form of storage. It acts as a cache to disk as it is much faster. A high utilization is better, as it indicates that more data is being cached. The ideal situation is when ESXi host Consumed metric is red but ESXi host Ballooned metric is green. When Ballooned is red and Consumed is gray, it means that there was high pressure in the past but it is not there anymore. The reason the ballooned stays red is because the ballooned pages were never requested back.

The ballooned memory counter was selected over the swapped or compressed memory counters as it is a better leading indicator. Since all three can co-exist at the same time, they are displayed in the line chart. Ballooned is displayed in absolute amount and not as a percentage, because the higher the size, higher are the chances for it to impact a VM. If you feel using percentage is easier for your operations, create a super metric to translate the value.

The heat map displays Wastage by a new color. The dark gray color indicates that wastage as capacity is not used. The performance problem due to low utilization can be caused by a bottleneck elsewhere.

Analyze if the ESXi host is contributing. A light gray box indicates that the host is a part of the cluster but there is no utilization. It is possible for the host to be in the maintenance mode or is powered off.

**Points to Note**

- ESXi host chooses to swap over compression if the compression ratio is less than 4x.
- If the ESXi host's physical NIC is saturated in your environment, then you can add a Network Throughput heat map.

### Live! Heavy Hitters Dashboard

The **Live! Heavy Hitters** dashboard helps you analyze the misuse of the shared infrastructure. This dashboard displays details of VMs misusing shared infrastructure and if that has caused performance problems to the other VMs. The shared infrastructure includes risks. The cause for excessive load might be attacks, for example, the denial of service, process runaway, or a mass activation of agents. The most demanding VM is the largest. If a handful of VMs is dominating the shared infrastructure, their collective size is displayed on the dashboard.

**Design Considerations**

See the Performance Dashboards page for common design considerations among all the dashboards for performance management.

In a shared environment, it is possible to have a victim-villain problem. In the heat map, the villain VM is the one with the largest box size, while the victim VM is the one with the red box. If a handful of VMs is dominating the shared infrastructure, their collective size will be highly visible on the dashboard.

**How to Use the Dashboard**

- The heat maps, Disk IOPS, Disk Throughput, Network Throughput, and CPU Demand displays the four different loads that can be excessive. The heat maps display the relative value and not the absolute value. A VM does not generate a high load in the absolute term just because it has a large configuration.
- Each heat map has its color threshold, reflecting the nature of the contention metrics used in each of them.
- For NOC Operator, drill-down by selecting one of the VMs on the heat map. All the four line charts are automatically displayed, enabling you to get a complete picture of the selected VM.

**Points to Note**

- Memory is not displayed as it is a form of storage. The memory counters are space utilization and not speed. Think of disk space instead of IOPS. It can cause a capacity problem on the shared ESXi host, but not performance problems to other VMs.
- In a large environment, it might be difficult to view a small victim VM. Consider having multiple dashboards and use them interchangeably.

### Live! vSphere VM Changes Dashboard

The **Live! vSphere VM Changes** dashboard displays changes to vSphere VMs as they happen. You can use the dashboard to ensure both the numbers and patterns are as per your expectation.

**How to Use the Dashboard**

The **Overview** widget at the top of the dashboard provides summary and overall high level vSphere VM numbers. Ensure that the numbers match your expectation. You can set thresholds so that it is easier to see if the values change (drop or increase) to a level that you are comfortable with. By default, you can view details of the last 6 hours. This allows you to view trends and patterns in the changes.

There are 3 types of detailed information covering different types of changes. This covers all the possible changes that could happen to a VM, hence providing real time visibility into the movement or volatility. The changes are categorized into

three groups for ease of understanding: Inventory, State, and Location. The changes in each widget are displayed in order of importance.

- **Inventory Changes:** The sum of the changes in VM inventory over the last 24 hours. Information is displayed about whether VMs are added or removed as per your expectation, if they are added or removed the right way (cloning vs free style creation vs template based deployment), and why there is a high count of VMs being unregistered or deleted.
- **Location Changes:** Refers to VMs that are moved to another host or datastore. This can be hot or cold migration. Hot migration is displayed first as it might impact performance because it changes both compute and storage. Pay attention to the storage migrations as they take the longest amount of time. If it takes longer than expected, something could be wrong. On the other hand, cold migration is typically an activity that should match the change request as the VMs are shutdown.
- **State Changes:** As reset is the least desired it is displayed first. It is used as the last resort. Suspend and Power Off are least preferred as these actions should be done from within the Guest OS. Expect these numbers to be low. A high number of powered on VMs could lead to high demand.

There are two types of widgets for each of the 3 types of changes:

- The first type of widget is a trend line. In this case the widgets are **Changes in VM Inventory**, **Changes in VM Location**, and **Changes in VM State**. Ensure that both the absolute amount and the pattern match your expectation.
- The second type of widget is a scoreboard showing the present number for each of the changes. In this case the widgets are **Details of VM Inventory Changes**, **Details of VM Location Changes**, and **Details of VM State Changes**. The change is sorted from the least desired (most impact to operations) to the least important.

**Points to Note**

The changes are color coded. You can adjust the thresholds for each change, however it is recommended that you avoid too many variations as it can get confusing. You can also adjust the widget size as required for large displays.

## Software Defined Wide Area Network Dashboard

The Software-Defined Wide Area Network (SD-WAN) dashboard allows you to configure and monitor the services related to VeloCloud and SD-WAN using VMware Aria OperationsVMware Cloud Foundation Operations. Using the SD-WAN dashboard, you can also collect the metrics for VeloCloud Orchestrator and VeloCloud Gateway.

By default the SD-WAN dashboards are deactivated, if you want to know how to activate them, see Manage Dashboards. The following services are discovered using the VeloCloud Orchestrator:

- Java Application
- VeloCloud Orchestrator
- Nginx
- ClickHouse
- MySQL
- Redis
- Network Time Protocol

The following services are discovered using VeloCloud Gateway:

- Network Time Protocol
- VeloCloud Gateway

### Troubleshoot SD-WAN Dashboard

You can use the widgets in Troubleshoot SD-WAN dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN.

You can use the dashboard widget in several ways:

- **Troubleshoot Virtual Machine (VM)**: Use this widget to navigate to a specific VM and troubleshoot the issues.
- **Troubleshoot Orchestrator**: Use this widget to navigate to a specific orchestrator and troubleshoot the issues.

- **Troubleshoot Gateway**: Use this widget to navigate to a specific gateway and troubleshoot the issues.
- **Troubleshoot Application**: Use this widget to navigate to a specific application and troubleshoot the issues.
- **Relationship**: Use this widget to view the services and operating system associated with the VeloCloud Orchestrator.
- **Top Alerts**: Use this widget to view the top alerts associated with the SD-WAN.

#### Troubleshoot SD-WAN Gateway Dashboard

You can use the widgets in Troubleshoot SD-WAN Gateway dashboard to monitor and troubleshoot all the services and applications associated with the SD-WAN gateway.

You can use the dashboard widget in several ways:
- **Active Alerts on the Gateway**: Use this widget to view the active alerts for the gateway.
- **Health of Gateway Applications**: Use this widget to view the health status of the applications in the gateway.
- **Examine Operating System**: Use this widget to examine the operating system status.
- **Gateway Summary Status**: Use this widget to view the summary information for the gateway.
- **Gateway Process Status**: Use this widget to view the process information for the gateway.
- **Gateway Resource Metrics**: Use this widget to view the resource metrics associated with the gateway.
- **Parent Host**: Use this widget to view the parent host information.
- **Parent Cluster**: Use this widget to view the parent cluster information.

#### Troubleshoot SD-WAN Orchestrator Dashboard

You can use the widgets in Troubleshoot SD-WAN Orchestrator dashboard to monitor and troubleshoot the services and applications associated with the SD-WAN Orchestrator.

You can use the dashboard widget in several ways:
- **Active Alerts on the Orchestrator**: Use this widget to view the active alerts for the Orchestrator.
- **Health of Orchestrator Applications**: Use this widget to view the health status of the applications in the gateway.
- **Examine Operating System**: Use this widget to examine the operating system status.
- **Examing MySQL**: Use this widget to examine the MySQL application.
- **Orchestrator Service Status**: Use this widget to view the service status of the Orchestrator.
- **Redis Status**: Use this widget to view the status of the Redis application.
- **API Check Status**: Use this widget to check the API status.
- **Nginx Status**: Use this widget to check the Nginx status.
- **Parent Host**: Use this widget to view the parent host information.
- **Parent Cluster**: Use this widget to view the parent cluster information.

## VMware Aria Automation Dashboards

The VMware Aria Automation dashboards allow you to track performance, health, utilization, and availability attributes of deployments made to SDDC clouds and public cloud end points. You can also monitor the virtual machines and track the utilization and performance of the VMs in your SDDCs and public cloud accounts..

#### VMware Aria Automation Predefined Dashboards

The following VMware Aria Automation dashboards are added to the predefined VMware Aria OperationsVMware Cloud Foundation Operations dashboards:
- Cloud Automation Environment Overview
- Cloud Automation Project Cost Overview
- Cloud Automation Resource Consumption Overview
- Cloud Automation Deployment Overview
- Cloud Automation Top-N Dashboard

### Automation Environment Overview

You can use the widgets in the Cloud Automation Environment Overview dashboard to view the environment details for the vCenter Cloud Zone objects. You can use the Cloud Automation Environment Overview dashboard to view the projects, deployments associated with the vCenter Cloud accounts.

You can use the dashboard widgets in several ways.

- **Environment Summary**: Use this widget to view the SDDC cloud zones, public cloud zones, projects, deployments, blueprints, SDDC VMs and public cloud VM details for the cloud accounts present in your environment.
- **SDDC Cloud Zones**: Use this widget to view the CPU, Disk, Memory, health, risk, and efficiency details for the SDDC cloud zone objects present in your environment.
- **Public Cloud Zones**: Use this widget to view the CPU, Disk, Memory, health, risk, and efficiency details for the public cloud zone objects present in your environment.
- **Project List**: Use this widget to view the total blueprints, cloud zones, deployments, virtual machines, health, risk, efficiency details in your environment.
- **Top Alerts**: Use this widget to view the top alerts in your environment.
- **VM List**: Use this widget to view all the VM details in your environment.
- **Blueprint List**: Use this widget to view the blueprint objects in your environment.
- **Deployment List**: Use this widget to view the blueprint objects deployed in your environment.
- **SDDC Virtual Machines**: Use this widget to view the SDDC VM's resource details.
- **Public Cloud Resources**: Use this widget to view the public cloud resource details.

### Automation SDDC Project Price Overview

You can use the cloud automation SDDC project price overview dashboard to view the project price details associated with each of VMware Cloud on AWS SDDC and public cloud accounts. With the project price overview dashboard, you can view price details for individual projects and find out deployments with the highest cost. .

You can use the dashboard widgets in several ways.

- **Project Cost**: Use this widget to view the project wise cost for compute, storage, and additional resources associated with your cloud environment.
- **Daily Price Over Time**: Use this widget to view the price of individual projects on a day to day basis.
- **Relationship**: Use this widget to view the relationship between the objects and the projects present in your cloud environment.
- **Deployment Price by Selected Project**: Use this widget to view the deployment cost for the selected project in your cloud environment.
- **Deployment With Highest Cost**: Use this widget to view the highest cost associated with projects that are present in your cloud environment.

### Automation SDDC Resource Consumption Overview

You can use the widgets in the Cloud Automation SDDC Resource Consumption Overview dashboard to view the resources consumed by VMware Aria Automation on Cloud Accounts.

You can use the Cloud Automation Resource Consumption Overview dashboard widgets in several ways.

- **Cloud Account**: Use this widget to view all the attributes related to the cloud account.
- **SDDC Cloud Zone**: Use this widget to view all the attributes related to the SDDC cloud zones.
- **Project**: Use this widget to view all the project details associated with your cloud account.
- **Cluster List**: Use this widget to view all the details associated with the clusters in your account.
- **Cluster Utilization**: Use this widget to view the cluster utilization details for the cloud accounts.
- **Deployment Heat Map by Project**: Use this widget to view the heat map for each deployed project in your cloud environment.
- **SDDC Cloud Zone Capacity**: Use this widget to view the memory and storage capacity that is allocated, reserved, and free for each cloud zone object.

- **SDDC Cloud Zone Memory Trend**: Use this widget to view and analyze a seven-day trend for the memory allocated, reserved, and free for the cloud zone.
- **SDDC Cloud Zone Storage Trend**: Use this widget to view and analyze a seven-day trend for the storage allocated, reserved, and free for the cloud zone.

### Automation Deployment Overview

You can use the cloud automation deployment overview dashboard to view the deployment details associated with your cloud environment. You can view details about cloud accounts, cloud zones, projects, and deployments. The dashboard also provides details about the deployment resources and the relationship between deployments and their objects.

You can use the dashboard widgets in several ways.
- **Cloud Account**: Use this widget to know the cloud account details of the selected account.
- **Cloud Zone**: Use this widget to know the cloud zone details like adapter, collection state, policy and object type of the selected account.
- **Project**: Use this widget to know the project details of the selected project.
- **Deployments**: Use this widget to know the deployment details of the cloud accounts in your environment.
- **Deployment Resources**: Use this widget to know the deployment resource details across the cloud accounts in your environment.
- **Object Relationship (Advanced)**: Use this widget to view the relationship between the objects and deployments present in your cloud environment.
- **Deployment Heat Map by Project**: Use this widget to view the heat map for each deployed project in your cloud environment.
- **Deployment Heat Map by Blueprint**: Use this widget to view the heat map for each deployed blueprint in your cloud environment.

### Automation Top-N Dashboard

You can use the widgets in the Cloud Automation Top-N dashboard to view the projects with most critical alerts, to view the blueprint with most deployments, and to view the deployments with the highest cost.

You can use the dashboard widgets in several ways.
- **Project with Most Critical Alerts**: Use this widget to view the projects which has most critical alerts.
- **Top Alerts**: Use this widget to view the top alerts for the projects in your cloud account.
- **Blueprints with Most Deployment**: Use this widget to view the blueprint which has maximum deployments for the cloud account.
- **Relationship**: Use this widget to analyze the relationship between blueprints and deployments, and deployment and cost.

### Project Chargeback

The Project Chargeback dashboard lets you know how much you must spend to run an VMware Aria Automation managed VM on behalf of your customer. In VMware Aria OperationsVMware Cloud Foundation Operations, you can configure the cost drivers and let the system automatically determine how much a VM costs based on your infrastructure requirement. Cost Drivers cover server hardware, storage, licenses, application, maintenance, labor, network, facilities, and additional costs configured within VMware Aria OperationsVMware Cloud Foundation Operations.

Price is what you charge your customer for running their VM. The price of a VM can be based on the cost of the project or based on a rate card that you define. Prices can include up charges, service charges, and others.

#### How to Use the Dashboard

- Select a Project widget displays the price of the group. The Select a Deployument widget lets you select deployments under the project, followed by the chargeback metrics.
- Price Summary of Selected Project shows the month to date price of the group.

- VM Price Distribution (Top 100) shows the most expensive VMs in the group.
- Powered Off VMs shows reclaimable VMs and their potential savings.
- Idle VMs shows reclaimable VMs and their potential savings.
- Reclaimable VMs with Snapshots shows reclaimable snapshots and their age.
- Price of VMs in the Selected Deployment shows the price and configuration of each VM in the selected group.

**Project Cost vs. Price**

The VM Cost vs. Price Dashboard helps you to analyze the relationship between cost and price for VMs managed by VMware Aria Automation. You can use this dashboard to ensure the price of VMs for chargeback is sufficient to cover the cost of running virtual machines.

**How to Use the Dashboard**

- Select a Project allows selection of a group of VMs to analyze. The Select a Deployment widget allows you to select deployments under the project, followed by the price/cost metrics.
- Summary (Month to Date) shows the month to date price and cost.
- VMs of the Selected Deployment (Select to View Trend) shows all VMs in the selected group with their Month to Date Cost, Today's Cost, Month to Date Price, and Today's Price.
- Daily Cost and Daily Price trend chart shows both cost and price over time.

**Project Showback**

The Showback VM Cost dashboard provides a quick Showback of the cost associated with the VMware Aria Automation managed VMs in a group. The definition of a group is based on several constructs such as Virtual Machine Folder, Custom Groups, VMware Aria Automation Projects, and more. Based on the Showback you can improve the accuracy of the costs by editing the cost drivers. Cost drivers that are not customized use reference cost, cost driver customization is available only in Advanced or Enterprise edition of VMware Aria Operations .

**How to Use the Dashboard**

- Select a Project widget displays the price of the group. The Select a Deployument widget lets you select deployments under the project, followed by the showback metrics.
- Cost Summary (This Month) shows the month to date cost, potential savings, and projected cost of the group.
- VM Cost Distribution (Top 100) shows the most expensive VMs in the group.
- Potential Savings (Top 10) shows the VMs ranked by their potential savings.
- VMs of the Selected Deployment (Select to View Trend) shows the cost and configuration of each VM in the selected group.
- Cost Trend of Selected VM shows the trend of the VMs cost over time.

## VMware Aria Operations Dashboards

With VMware Aria Operations dashboards you can monitor and troubleshoot objects in your cloud infrastructure.

The following v VMware Aria Operations dashboards are added to the predefined VMware Aria OperationsVMware Cloud Foundation Operations dashboards:

- VMware Aria Operations Cloud Billing
- VMware Aria Operations Cloud Universal Billing

**VMware Aria Operations Cloud Billing**

The VMware Aria Operations Cloud Billing dashboard provides you object billing details of Operating System Instance (OSI) used in your cloud environment.

**How to Use the Dashboard**

- OSIs and Billable Objects widget provides the total count of OSIs and billable objects. You have to update these widgets depending on your subscription limits.
- OSIs Across Object Types widget provides distribution of OSIs across different object types.
- Billable Object Types List widget provides a list of all the object types which are managed by VMware Aria Operations and consume license unit(s).
- OSIs Consumption Across Object Types displays a heat map and maps the magnitude of OSIs consumption for different object types with the relevant heat map colors. The widget also interacts with OSIs Consumption Over Time and displays how OSIs count have been changing for a given object type over a period of time.
- List of Objects widget displays the object details like name, adapter type, object type, policy, creation time, collection state, and collection status. Use the filter option to filter different objects from the list of objects.

**How to Edit OSIs and Billable Objects Widget To Set Correct Color Codes**

1. Click the Edit icon at the top right corner of the widget.
2. Go to **Output Data** section.
3. Double click the row which has column heading Yellow, Orange, or Red.
4. Set the **Color Method** to custom.
5. Enter appropriate values as per your subscription limits.

**VMware Aria Operations Cloud Universal Billing**

The VMware Aria Operations Cloud Universal Billing dashboard provides the object billing details based on CPU usage.

**How to Use the Dashboard**

- CPUs and Billable Objects widgets provides the total count of CPUs and Billable Objects. You have to update these widgets based upon your subscription limits.
- CPUs Across Object Types widget provides distribution of CPUs across different object types.
- Billable Object Types List widget provides a list of all the object types which are actively managed by VMware Aria Operations and consume license unit(s).
- CPUs Consumption Across Object Types displays a heatmap and maps the magnitude of CPUs consumption for different object types with heatmap colors. The widget also interacts with CPUs Consumption Over Time and displays how CPUs count has been changing for a given object type over a period of time.

**How to Edit CPUs and Billable Objects Widget To Set Correct Color Codes**

1. Click the Edit icon at the top right corner of the widget.
2. Go to **Output Data** section.
3. Double click the row which has column heading Yellow, Orange, or Red.
4. Set the **Color Method** to custom.
5. Enter appropriate values as per your subscription limits.

## Inventory Dashboards

The three vSphere Inventory dashboards and workload management inventory dashboards cater to the compute, network, and storage aspects of your SDDC. Using these dashboards, you can navigate through the environment and view your inventory and their key metrics at a glance. The Network and Storage dashboards can be shared with the network and storage teams respectively, giving them the necessary visibility, and increasing the collaboration between teams.

## vSphere inventory dashboards

The vSphere inventory dashboards are built specifically for each role, but they share a common design. They have a similar layout and are used in the same manner. This makes learning easier, especially in smaller environments where the same team manages the full environment.

These dashboards help you answer several key questions:
- What is the topology of your vSphere compute inventory?
- What is the topology of your vSphere storage inventory?
- What is the topology of your vSphere network inventory?

## Workload Management Inventory Dashboard

This is a unified dashboard for the new workload management objects. It shows the relationships and KPIs for the workload management objects. For example, you can see the topology view from the Tanzu Kubernetes clusters to the physical infrastructure.

### vSphere Compute Inventory Dashboard

You can use the vSphere Compute Inventory Dashboard to browse through the topology of your vSphere compute inventory which includes information related to vSphere world, vCenter, data center, clusters, hosts, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the clusters, ESXi hosts, and virtual machines associated with the object.

You can use the dashboard widgets in several ways.
- **Properties**: View the properties related to an object in the environment.
- **Metrics**: View the metrics related to the object.
- **Clusters**: View the cluster functionality.
- **ESXi Hosts**: View the data related to the hosts.
- **Virtual Machines**: View VMs that belong to the object.

### vSphere Network Inventory Dashboard

The vSphere Network Inventory Dashboard allows you to browse through the topology of your vSphere network inventory which includes information related to vSphere world, vCenter, data center, distributed vSwitches, distributed port groups, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the distributed vSwitches, distributed port groups, virtual machines associated with it.

You can use the dashboard widgets in several ways.
- **Properties**: View the properties related to the object in the environment.
- **Metrics**: View the metrics of the object.
- **Distributed vSwitches**: View details related to the distributed vSwitches.
- **Distributed Port Groups**: View data relevant to distributed port groups.
- **Virtual Machines**: View VMs that belong to the object.

### vSphere Storage Inventory Dashboard

The vSphere Storage Inventory dashboard allows you to browse through the topology of your vSphere storage inventory which includes information related to vSphere world, vCenter, data center, datastore clusters, datastores, virtual machines, properties, and metrics.

You can select an object type to view the properties and metrics related to it. You can also view the datastore clusters, datastores, and virtual machines associated with it.

You can use the dashboard widgets in several ways.
- **Properties**: View the properties related to the object in the environment.
- **Metrics**: View the metrics of the object.
- **Datastore Clusters**: View the datastore cluster functionality.
- **Datastores**: View the datastore functionality.
- **Virtual Machines**: View VMs that belong to the object.

#### Workload Management Inventory Dashboard

The Workload Management Inventory dashboards curates the Kubernetes inventory across all the Workload Management activated vSphere environments and displays it here. This includes an end to end topology map showcasing the health of all the objectes along with upstream and downstream dependencies. Upon clicking any object in the relationship tree, the related inventorty of Supervisor Clusters, Namespaces, Pods, Developer Managed VMs and Tanzu Kubernetes clusters can be viewed and exported from this dashboard.

You can select an object type to view the properties and key metrics related to it.

You can use the dashboard widgets in several ways.
- **Environment Summary**: Provides a summary of the supervisor cluster and the child objects.
- **Relationships**: An interactive canvas where you can view the relationship between the different objects in the workload management inventory.
- **Properties**: View the properties related to the object in the environment.
- **Metrics**: View the metrics of the object.
- **Supervisor Clusters**: View the supervisor cluster functionality.
- **Tanzu Kubernetes cluster**: View the Tanzu Kubernetes cluster functionality.
- **Virtual Machines**: View VMs that belong to the object.
- **vSphere Pods**: View information about vSphere Pods.

#### vSphere VM Inventory Dashboard

The vSphere VM Inventory dashboard provides real-time updates on changes to virtual machines, allowing users to monitor both the quantity and patterns of these changes to ensure they align with expectations. To facilitate understanding, the 20 types of changes are categorized into three groups, with changes within each group ordered by importance. For instance, deletions are prioritized over deployments. This tool is specifically designed for the Network Operations Center support team to help them efficiently track and manage VM modifications.

##### The vSphere VM Inventory Dashboard Display

The Overview widget provides high-level metrics, displaying data from the last six hours by default to help identify patterns. The Changes in VM Location section tracks movements across ESXi hosts and datastores, including vSAN. The Changes in VM States section monitors changes in power or running states, detailing how these changes were executed.

##### Customizing the vSphere VM Inventory Dashboard

Adjust the threshold settings to suit your specific needs, but be cautious of setting too many variations as it can become confusing. If you operate with two physical data centers, consider splitting the dashboard into two columns, one for each data center. For those using the dashboard on a desktop, adding a navigation link to the vSphere VM Inventory dashboard can enhance convenience.

## Dashboards in VMware Cloud on AWS

The **VMware Cloud on AWS** dashboards allow you to track the capacity, cost, and inventory overviews of the SDDCs. You can also track the virtual machines monitoring and the utilization and performance of these SDDCs.

### VMC Capacity Dashboard

Use the **VMC Capacity** dashboard to view the capacity overview of each VMware Cloud on AWS SDDC. You can view the capacity of Clusters, Hosts, VMs, Datastores, and Disk groups.

**Table 187: Widgets in VMC Capacity Dashboard**

| Widget | Description |
|---|---|
| VMC SDDC by Capacity Remaining % | Displays the SDDCs as cards that show the remaining capacity percentage. |
| VMC SDDC by Time Remaining % | Displays the SDDCs as cards that show the remaining time percentage. |
| VMC SDDC by Virtual Machine Remaining (based on avg VM profile) | Displays the SDDCs as cards that show the remaining number of virtual machines. |

When you select one of the SDDC cards, the details of that SDDC are automatically populated in the widgets after the VMC SDDC by Virtual Machine Remaining (based on avg VM profile) widget.

**NOTE**
The key kpis are color-coded to help in identifying capacity bottlenecks.

### VMC Cost Overview Dashboard

Use the **VMC Cost Overview** dashboard to view the organization cost overview and expense trends. The monthly metrics plotted in the trends represent the previous month's bill. The bill start date and end date are available in the properties.

**Table 188: Widgets in VMware Cloud on AWS Dashboard**

| Widget | Description |
|---|---|
| Organization Cost Overview | Displays a list of organizations with the details of their Outstanding Expense, Commit Expense (YTD), On Demand Expense (YTD), and Total Expense (YTD). |
| Outstanding Expense Trend | Displays the outstanding expense trend of the organization selected in the Organization Cost Overview widget. |
| Total Expense Trend (Monthly) | Displays the total monthly expenses trend of the organization selected in the Organization Cost Overview widget. |
| Commit Expense Trend (Monthly) | Displays the committed monthly expense trend of the organization selected in the Organization Cost Overview widget. |
| On-Demand Expense Trend (Monthly) | Displays the on-demand monthly expenses trend of the organization selected in the Organization Cost Overview widget. |
| Purchase History | Displays the bill line items/purchases from the available bills. |
| Currency Information | Represents the metrics currency unit set in this management pack account. |

**NOTE**
The YTD metric is an aggregation from the beginning of the calender year, until the last available bills.

### VMC Inventory Dashboard

Use the **VMC Inventory** dashboard to view the inventory overview of all the SDDCs configured in VMware Cloud on AWS.

**Widgets in VMC Inventory Dashboard**

VMC SDDCs: displays the SDDCs as cards that show the number of virtual machines running in the SDDC. The SDDC card also shows a trend of virtual machine growth over the past 30 days. If you are about to reach the limit of supported virtual machines in that SDDC, the SDDC card indicates this by changing colors.

When you select one of the SDDC cards, the list of all the vSphere Clusters, Datastores, vSphere Hosts, and VMs with key configuration details of that SDDC are populated in the widgets after the VMC SDDCs widget.

You can choose to export the desired list in a CSV format using the toolbars on the widget list.

### VMC Management VM Monitoring Dashboard

Use the **VMC Management VM Monitoring** dashboard to monitor the utilization and performance of the key management VMs running in your SDDC. This dashboard ensures that the management components (such as vCenter and NSX) are not facing any resource bottlenecks from the CPU, memory, network, and storage perspectives.

**Table 189: Widgets in VMC Management VM Monitoring Dashboard**

| Widget | Description |
|---|---|
| CPU Usage & Performance | Displays the list of all the management components in each SDDC with key CPU utilization and performance KPIs. Select a management VM to see the usage and performance trends of all the CPU cores. |
| Memory Usage & Performance | Displays the list of all the management components in each SDDC with key Memory utilization and performance KPIs. Select a management VM to see the memory usage and performance trends. |
| Network Usage & Performance | Displays the list of all the management components in each SDDC with key Network utilization and performance KPIs. Select a management VM to see the memory usage and performance trends. |
| Storage Usage & Performance | Displays the list of all the management components in each SDDC with key storage utilization and performance KPIs. Select a management VM to see the network usage and performance trends. |

### VMC Utilization and Performance Dashboard

Use the **VMC Utilization and Performance** dashboard to view the utilization and performance overview of each SDDC based on heavy hitter VMs and impacted VMs over the last 30 days. This dashboard helps you in finding the VMs in your environment that are negatively impacting the capacity or performance from a CPU, memory, storage, or network perspective.

**Widgets in VMC Utilization and Performance Dashboard**

List of VMC SDDCs: displays the list of all the SDDCs with aggregate CPU, memory, and storage utilization with 95th percentile and maximum values over the last 30 days.

When you select one of the SDDC from the List of VMC SDDCs widget, you can see the list of top VMs that are consuming compute, network & storage resources in that SDDC. The widgets after that show the compute (CPU & memory) utilization and performance analysis, network, storage, and utilization and performance analysis.

Each section in the dashboard is based on the last 30 days data with 95th percentile transformation which is configurable to Max, Average, Current, Standard Deviation, or other mathematical transformations.

### VMC Configuration Maximums Dashboard

Use the **VMC Configuration Maximums** dashboard to view the VMC limits and your consumption against those limits. This dashboard displays alerts for configuration maximum, and details of organization, SDDC, vSAN, and cluster maximums.

**Table 190: Widgets in VMC Configuration Maximums Dashboard**

| Widget | Description |
|---|---|
| Select an Environment | Select an environment for which you want to view the alerts and other details. Once you select an environment, the details of that environment are automatically populated in the widgets below. |
| VMC Configuration Maximums Alerts | Displays the list of alerts for the selected environment. |
| Number of SDDCs | Displays the number of SDDCs for the organization maximums, the provisioned, and the soft limit used. |
| Number of Hosts | Displays the number of hosts for the organization maximums, the provisioned, and the soft limit used. |
| Public IP Addresses (Elastic IPs) | Displays the public IP addresses for the organization maximums, the provisioned, and the soft limit used. |
| Maximum Clusters | Displays the maximum clusters for the SDDC maximums, the provisioned, and the hard and soft limit used. |
| Maximum Hosts | Displays the maximum hosts for the SDDC maximums, the provisioned, and the limit used. |
| Maximums VMs | Displays the maximum VMs for the SDDC maximums, the provisioned, and the limit used. |
| Linked VPCs | Displays the linked VPCs for the SDDC maximums, the provisioned, and the limit used. |
| Clusters with No SLA | Displays the maximum number of clusters and the number of provisioned clusters with no SLA per SDDC. An empty list means no clusters have been identified with no SLA. |
| Clusters with Limited SLA | Displays the maximum number of clusters and the number of provisioned clusters with limited SLA per SDDC. An empty list means no clusters have been identified with limited SLA. |
| Max hosts per Cluster (including stretched clusters) | Displays the maximum hosts per cluster including the stretched clusters, the provisioned, and the limit used. |
| Datastore Utilization | Displays the datastore utilization for vSAN maximums, the used space, utilization limit, and the remediation needed. |
| VMs per Host Limit Used | Displays the maximum number VMs that can be deployed per host, VMs that are provisioned per host, and the percentage of limit used. |
| VMs per Host Limit Used of Selected Host | Displays the VMs that are used per host limit for a selected host. |

## Google Cloud VMware Engine Dashboards

The Google Cloud VMware Engine dashboards allow you to track the configuration maximums, cost, and inventory overviews of the Google Cloud VMware Engine private clouds.

### GCVE Configuration Maximums Dashboard

Use the **GCVE Configuration Maximums** dashboard to view the Google Cloud VMware Engine limits and your consumption against those limits. This dashboard displays alerts for configuration maximum, and details of the adapter instance, private cloud, vSAN, and host and cluster maximums.

**Table 191: Widgets in the GCVE Configuration Maximums Dashboard**

| Widget | Description |
|---|---|
| Select an Environment | Select an environment for which you want to view the alerts and other details. Once you select an environment, the details of that environment are automatically populated in the widgets below. |
| GCVE Configuration Maximums Alerts | Displays the list of alerts for the selected environment. |
| **Private Cloud Maximums** | |
| Maximum Clusters | Displays the maximum clusters for Google Cloud VMware Engine private cloud maximums, the provisioned, and the hard and limit used. |
| Maximum Hosts | Displays the maximum hosts for Google Cloud VMware Engine private cloud maximums, the provisioned, and the hard limit used. |
| **Cluster Maximums** | |
| Max hosts per Cluster (including stretched clusters) | Displays the maximum hosts per cluster including the stretched clusters, the provisioned, and the limit used. |
| **vSAN Maximums** | |
| Datastore Utilization | Displays the datastore utilization for vSAN maximums, the used space, utilization limit, and the remediation needed. |

### GCVE Cost Overview Dashboard

Use the **GCVE Cost Overview** dashboard to view the organization cost overview and expense trends. The monthly metrics plotted in the trends represent the previous month's bill. The bill start date and end date are available in the properties.

**Table 192: Widgets in the GCVE Cost Overview Dashboard**

| Widget | Description |
|---|---|
| Organization Cost Overview | Displays a list of bills with the details of their Outstanding Expense, Commit Expense (YTD), On Demand Expense (YTD), and Total Expense (YTD). |
| Outstanding Expense Trend | Displays the outstanding expense trend of the bill selected in the **Organization Cost Overview** widget. |
| Total Expense Trend (Monthly) | Displays the total monthly expenses trend of the bill selected in the **Organization Cost Overview** widget. |
| Commit Expense Trend (Monthly) | Displays the committed monthly expense trend of the bill selected in the **Organization Cost Overview** widget. |
| On-Demand Expense Trend (Monthly) | Displays the on-demand monthly expenses trend of the bill selected in the **Organization Cost Overview** widget. |
| Purchase History | Displays the bill line items/purchases from the available bills. |
| Currency Information | Represents the metrics currency unit set in this management pack account. |

**NOTE**
The YTD metric is an aggregation from the beginning of the calender year, until the last available bills.

### GCVE Inventory Dashboard

Use the **GCVE Inventory** dashboard to view the inventory overview of all the private clouds configured in Google Cloud VMware Engine.

**Widgets in the GCVE Inventory Dashboard**

**GCVE Private Clouds**: Displays a card per Google Cloud VMware Engine private cloud with the number of virtual machines running in each private cloud. The card also displays a trend of virtual machine growth over the past 30 days. If you are about to reach the limit of supported virtual machines in that private cloud, the card indicates this by changing colors.

When you select one of the private cloud cards, the widgets below are populated and you can see the list of all the vSphere Clusters, Datastores, vSphere Hosts, and VMs in that private cloud with key configuration details.

You can choose to export the desired list in a CSV format using the toolbars on the widget list.

## Dashboards in NSX Management Pack

The **NSX Main** dashboard provides an overview of the network objects. It displays the topology of a selected object, how it connects to the elements in the network, and a view of related alerts.

**Table 193: Widgets in NSX Main Dashboard**

| Widget | Description |
|---|---|
| NSX Instances | Displays the list of environments that are being monitored. When you select an environment in this widget, the other widgets in the **NSX Main** dashboard display data for the selected adapter. |
| Environment Overview | Displays a top-level view of the selected environment and following key components.<br>• NSX Manager<br>• Controller Node<br>• Logical Router<br>• Logical Switch<br>• Load Balancer Virtual Server<br>• Transport Zone |
| Top Alerts | Displays all the open alerts for the selected object in the **Environment Overview** widget. |
| Topology Graph | Displays the topology of the selected object in the **Environment Overview** widget. |

### NSX Configmax Metrics

The **NSX Configmax Metrics** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations provides an overview of all the configuration maximum metrics in all the NSX instances.

**Table 194: Widgets in NSX Configmax Metrics Dashboard**

| Widget | Description |
|---|---|
| Select an adapter instance | Displays the list of all the NSX and NSX on VMC instances. When you select an instance in this widget, the other widgets in the **NSX Configmax Metrics** dashboard display data for the selected instance. |
| Relationship view | Displays the objects hierarchy for the instance selected in the **Select an adapter instance** widget. Only the objects with configuration maximum metrics are shown in the relationship view. |
| Select object from relationship view for the configmax metric | Displays all the configmax metrics for the selected object in the **Relationship View** widget. |
| Trend View | Displays all the MGW, CGW, and Distributed firewall section rule trends of the instance selected in the **Select an adapter instance** widget.<br>**NOTE**<br>The Trend View widget loads the trends only for the firewall sections object on VMware Cloud on AWS instances. |

**NSX Configmax Metrics**

The NSX Configmax Metrics provides information about the new metrics added to the NSX Configmax Metrics dashboard.

**Table 195: NSX Configmax Metrics**

| Metric Type | Metric | Description |
|---|---|---|
| Group | Configuration Maximums\|Count\|Tag Count | This metric displays the number of tags for the selected group. |
| Logical Router | Configuration Maximums\|ARP Entries Count | This metric displays the number of ARP entries of the logical router. |
| | Configuration Maximums\|Router Port Count | This metric displays the number of ports available in the logical router. |
| Management Cluster | Configuration Maximums\|Prepared vC Cluster Count | This metric displays the number of prepared vCenter clusters in the management cluster. |
| | Configuration Maximums\|Compute Manager Count | This metric displays the number of compute manager present in the management cluster. |
| Edge Cluster | Configuration Maximums\|Edge Node Count | This metric displays the number of edge nodes present in the edge cluster. |

**NSX Inventory Dashboard**

The **NSX Inventory** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations provides a holistic view of NSX environment topology. It offers an graphical representation of the interconnections between various NSX components.

This dashboard helps administrators with essential insights into the following components:

- NSX Management Nodes: Displays information related to the status and configuration of NSX management nodes.
- T0 and T1 Routers: Displays details of the t0 and t1 routers including their configurations and associations.
- NSX Load Balancers (LBs) and LB Pools: Displays information about NSX Load Balancers, virtual servers, pool members, and health checks.
- Edge Clusters: Displays data on Edge Clusters including cluster membership and resource consumption.
- Transport Zones: Displays insights into transport zones, offering visibility into network segmentation and connectivity options.

- Edge Nodes and Host Nodes: Monitors Edge Nodes and Host Nodes that are crucial components of the NSX infrastructure.
- Logical Switches: Displays a comprehensive list of logical switches with associated properties and configurations.
- NSX Groups: Displays details related to NSX Security Groups and Service Groups, simplifying group membership and policy management.

### NSX System Dashboard

The **NSX System** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations provides essential insights into the core components of the NSX infrastructure specific to inventory and health.

This dashboard offers visibility into the following aspects:
- Management Clusters: Displays information about NSX Management Clusters including cluster health, member nodes, and their roles.
- Advanced Object Relationship: Allows you to explore intricate connections between different NSX objects, enhancing understanding of component interactions.
- Summary: Displays a concise snapshot of NSX system, presenting key performance metrics and alert notifications.
- Alerts: Monitors and manages alerts generated within the NSX environment, offering insights into critical issues that require attention and resolution.

### NSX Edge Dashboard

The **NSX Edge** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations provides detailed insights into the performance and configuration of Logical Routers, which is a foundational component of the NSX Edge environment.

This dashboard equips administrators with necessary tools to effectively manage, optimize, and troubleshoot the performance of Logical Routers within the NSX Edge environment. This dashboard offers the following features:
- Logical Routers: Displays information related to Logical Routers, including configurations, routing tables, and interface details.
- Advanced Object Relationship for Logical Router: Displays insights into complex relationships specific to logical routers.
- Logical Router RX and TX: Monitors traffic received (RX) and traffic transmitted (TX) by logical routers to optimize network performance and troubleshoot issues.
- VMs Using Logical Router: Identifies virtual machines (VMs) utilizing specific logical routers, offering visibility into network traffic patterns and dependencies.

### NSX Switch Dashboard

The **NSX Switch** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations provides information about logical switches. Its includes an inventory of logical switches with their associated properties and configurations.

Additionally, the dashboard offers the following functionalities:
- Inventory of Logical Switches: Displays detailed information on logical switches for efficient tracking and management of network segmentation.
- List of NSX Logical Switches: Displays a comprehensive list of NSX logical switches, simplifying the tracking of network segmentation.
- Advanced Object Relationship of Logical Switch: Provides insights into the intricate relationships and dependencies between logical switches and other NSX components.
- Logical Switch Inbound and Outbound Throughput: Monitors inbound and outbound traffic throughput for logical switches to optimize network performance and identify potential bottlenecks.

### NSX Transport Node Performance Dashboard

The **NSX Transport Node Performance** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations provides an extensive view of performance metrics related to NSX Transport Nodes. It provides essential insights into the operation and efficiency of these nodes, aiding network administrators in optimizing resource allocation and ensuring seamless data packet forwarding within the NSX environment.

You can monitor the following key metrics:
- DPDK CPU Core Average Usage
- Non-DPDK CPU Core Average Usage
- Memory used for Transport Node
- Transport Node Interface Statistics

### NSX Edge Performance Dashboard

The **NSX Edge Performance** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations provides performance information on NSX Edge devices. It offers visibility into the operation of these devices, assisting users in maintaining optimal routing and network services while swiftly addressing potential issues.

You can monitor the following key metrics:
- List of NSX Logical Routes
- RX MBPs (Megabits per second)
- TX MBPS (Megabits per second)
- RX Packets Per Second
- TX Packets Per Second
- RX Packets Dropped Per Second
- TX Packets Dropped Per Second

### NSX Switch Performance Dashboard

The **NSX Switch Performance** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations focuses on the performance of logical switches of the NSX environment. It simplifies the tracking and management of network segmentation while offering insights into the dependencies between logical switches and other NSX components.

You can monitor the following key metrics:
- List of NSX Logical Switches
- Advanced Object Relationships
- Performance Statistics (Inbound and Outbound Throughput)

### NSX Load Balancer Performance Dashboard

The **NSX Load Balance Performance** dashboard in VMware Aria OperationsVMware Cloud Foundation Operations offers detailed insights into the performance of NSX Load Balancer services. It includes critical data related to the operation of Load Balancers, providing administrators with the necessary information to optimize configurations and maintain high availability and performance for applications and services.

You can monitor the following key metrics:
- CPU Usage of the Load Balancer
- Memory Usage of the Load Balancer
- Session Statistics

## Aggregator Management PackCloud Federation Adapter Dashboards

Aggregator Management PackCloud Federation Adapter caters to certain out-of- the-box use cases, which are delivered through dashboards. After the installation and configuration of the Cloud Federation Adapter management pack, you can access these out-of- the-box dashboards.

Ensure that you have enabled the following Management Packs for the Cloud Federation Adapter dashboards to display data.

- VMware Aria Automation
- CloudHealth
- VMware Cloud on AWS
- AWS
- Azure

Verify that you have installed the Aggregator Management Pack.

1. To access the dashboards, in the left pane of the VMware Aria OperationsVMware Cloud Foundation Operations, click **Dashboards**.
2. Click **AggregatorCloud Federation Analytics** to view all dashboards.
3. From the dashboard list, select **Getting Started- AggregatorGetting Started - Cloud Federation** dashboard in the list of all the dashboards and associated use cases.

### Getting Started - Aggregator Getting Started - Cloud Federation Dashboard

The Getting Started - AggregatorGetting Started - Cloud Federation dashboard displays the catalog of AggregatorCloud Federation Analytics dashboards and the AggregatorCloud Federation Analytics deployment map.

You can step through the multiple categories under AggregatorCloud Federation Analytics page to cater to specific use cases. The deployment map on this page provides a topology of your AggregatorCloud Federation Analytics Deployment with colors depicting the current health and collection status.

**Figure 11: Getting Started - Aggregator Dashboard**

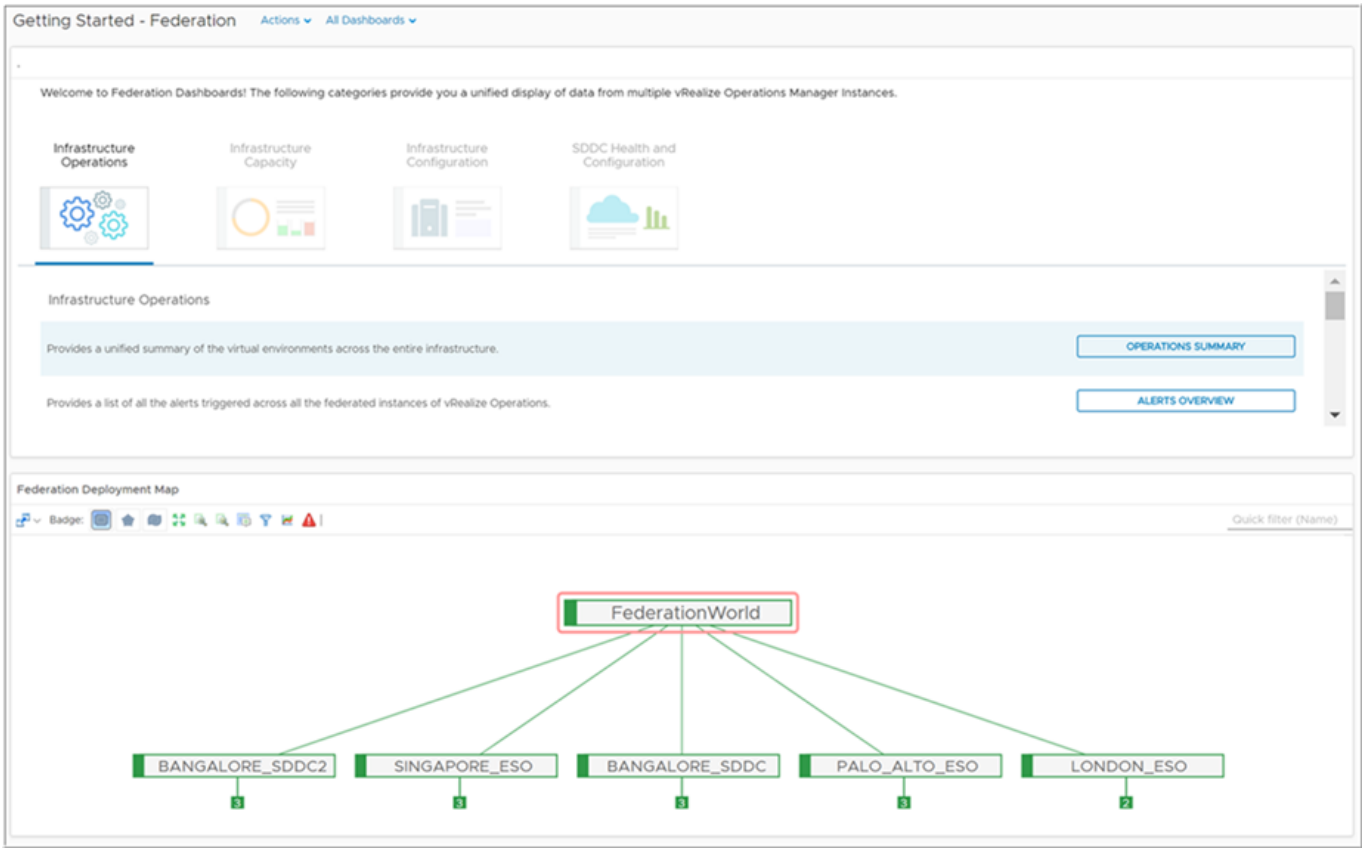

## Infrastructure Operations

The Getting Started page provides access to the following OOTB categories. The infrastructure category caters to Senior Executives and NOC users by providing them an operations summary and alerts. The dashboards under this category are.

- Operations Summary
- Alerts Overview

### Operations Summary Dashboard

The Operations Summary dashboard provides a count of your infrastructure inventory objects across all the VMware Aria OperationsVMware Cloud Foundation Operations instances being monitored by the vRealize Operations Federation Management PackCloud Federation Adapter. It provides you the overall alert volume in these environment, along with a summarized view of health and capacity of the vSphere clusters in your environment.

**Figure 12: Operations Summary Dashboard**

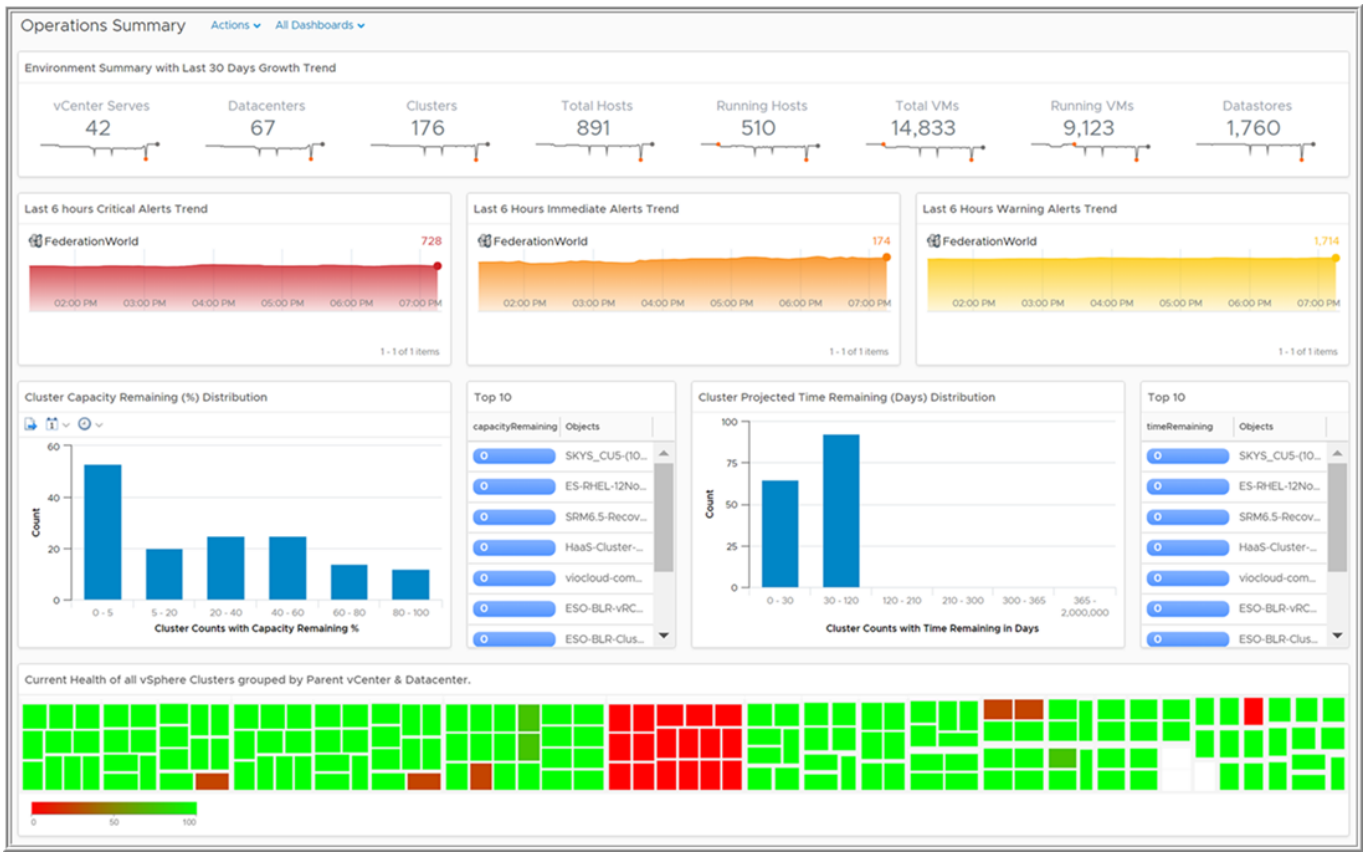


### Alerts Overview

The Alerts Overview dashboard provides the alerts page of VMware Aria OperationsVMware Cloud Foundation Operations. On this page, you can not only view the alerts triggered on your VMware Aria OperationsVMware Cloud Foundation OperationsAggregatorCloud Federation Analytics instance, but also view the alerts triggered in your child VMware Aria OperationsVMware Cloud Foundation Operations instances.

The overview page includes all the alerts triggered on the VMware Aria OperationsVMware Cloud Foundation Operations instance and is not dependent on the objects you are monitoring from the child VMware Aria OperationsVMware Cloud Foundation Operations instance.

This page can help your NOC teams to have a single pane for all the alerts across your environments. To view the details of a specific alert, click on that alert to get directed to the child VMware Aria OperationsVMware Cloud Foundation Operations instance where the alert was originally triggered for troubleshooting purposes.

**Figure 13: Alerts Overview**

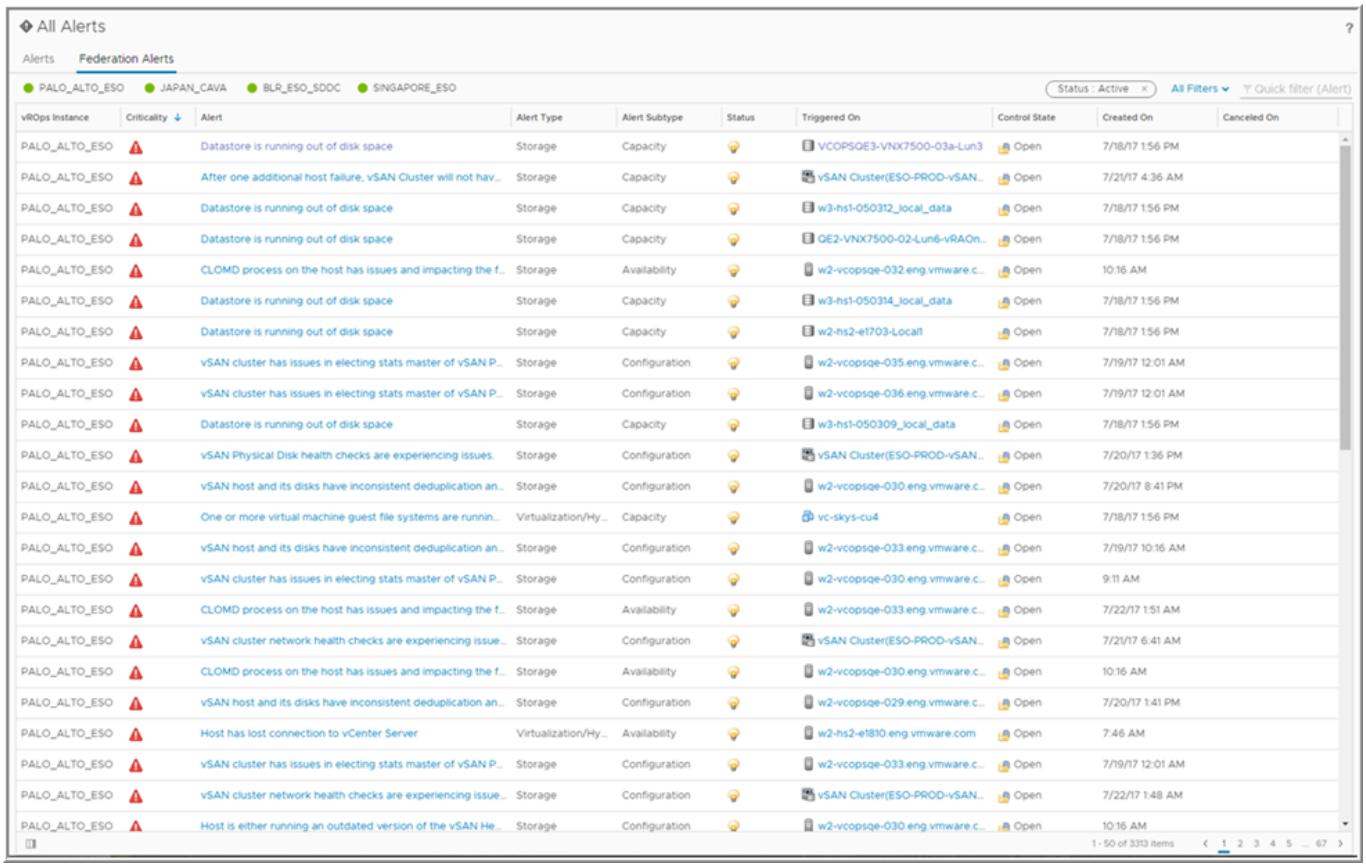


## Infrastructure Capacity

The Infrastructure Capacity category caters to Senior Executives and NOC users by providing them an overview of overall capacity and reclamation opportunities.

The following dashboards are under the Infrastructure Capacity:

- Capacity Summary
- Capacity Optimization

## Capacity Summary Dashboard

The Capacity Summary Dashboard provides you with a summary of the total physical capacity available across all your environments being monitored by multiple instances of VMware Aria OperationsVMware Cloud Foundation Operations. The dashboard also provides you a list view of all your clusters across your environment with details around inventory, capacity and utilization of those clusters. You can export this list into a CSV file for reporting purposes.

**Figure 14: Capacity Summary Dashboard**

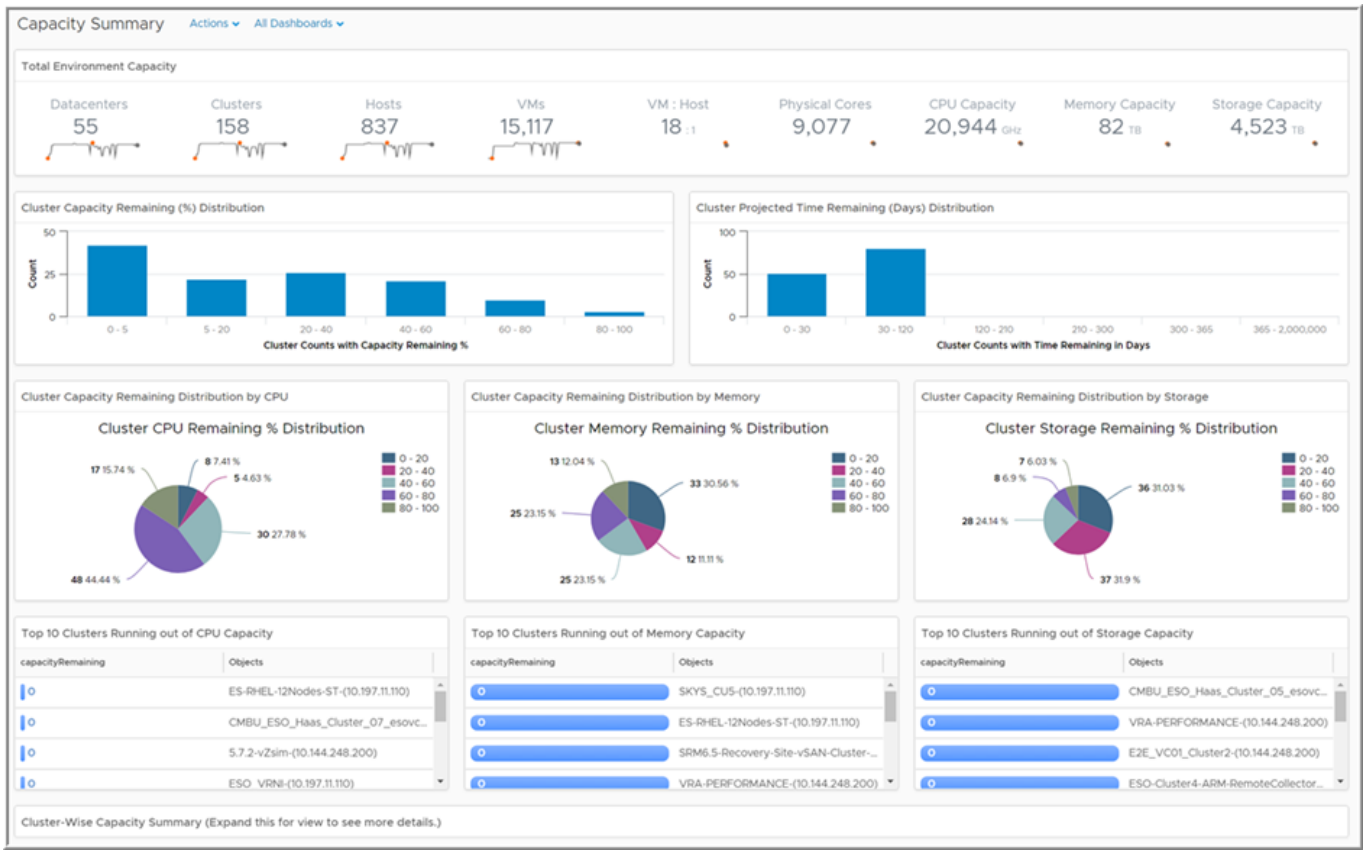


## Capacity Optimization Dashboard

The Capacity Optimization dashboard provides you a quick view of resource optimization opportunities within your virtual infrastructure. This dashboard is focused on improving the efficiency of your multiple instances of VMware Aria OperationsVMware Cloud Foundation Operations environment by reducing the wastage of resources.

**Figure 15: Capacity Optimization Dashboard**

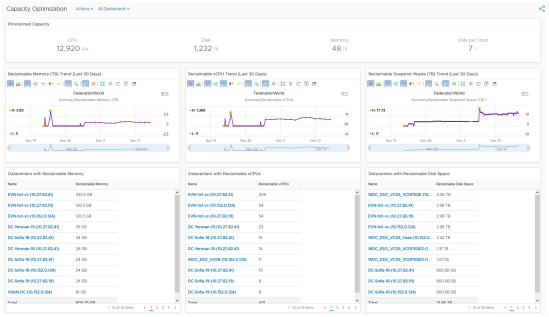

## Infrastructure Configuration

Infrastructure Configuration category caters to Senior Executives and virtual infrastructure admin by providing them the summary of ESXi Host and cluster configurations.

The dashboards under Infrastructure Configuration category are:

- vSphere Cluster Configuration
- vSphere Host Configuration

### vSphere Cluster Configuration Dashboard

The vSphere Cluster Configuration dashboard provides a quick overview of your vSphere cluster configurations from various instances of VMware Aria OperationsVMware Cloud Foundation Operations. The dashboard highlights the areas that are important in delivering performance and availability to your virtual machines.

**Figure 16: vSphere Cluster Configuration Dashboard**


The dashboard also highlights if there are vSphere clusters which are not configured for DRS, High Availability (HA), or admission control to avoid any resource bottlenecks or availability issues when a host fails. The heat map in this dashboard helps you to identify if you have hosts where vMotion was not enabled as this may not allow the VMs to move from or to that host. This can cause potential performance issues for the VMs on that host if the host gets too busy.

### vSphere Host Configuration Dashboard

The vSphere Host Configuration dashboard provides an overview of your ESXi host configurations, and displays inconsistencies so that you can take corrective action.

**Figure 17: vSphere Host Configuration Dashboard**

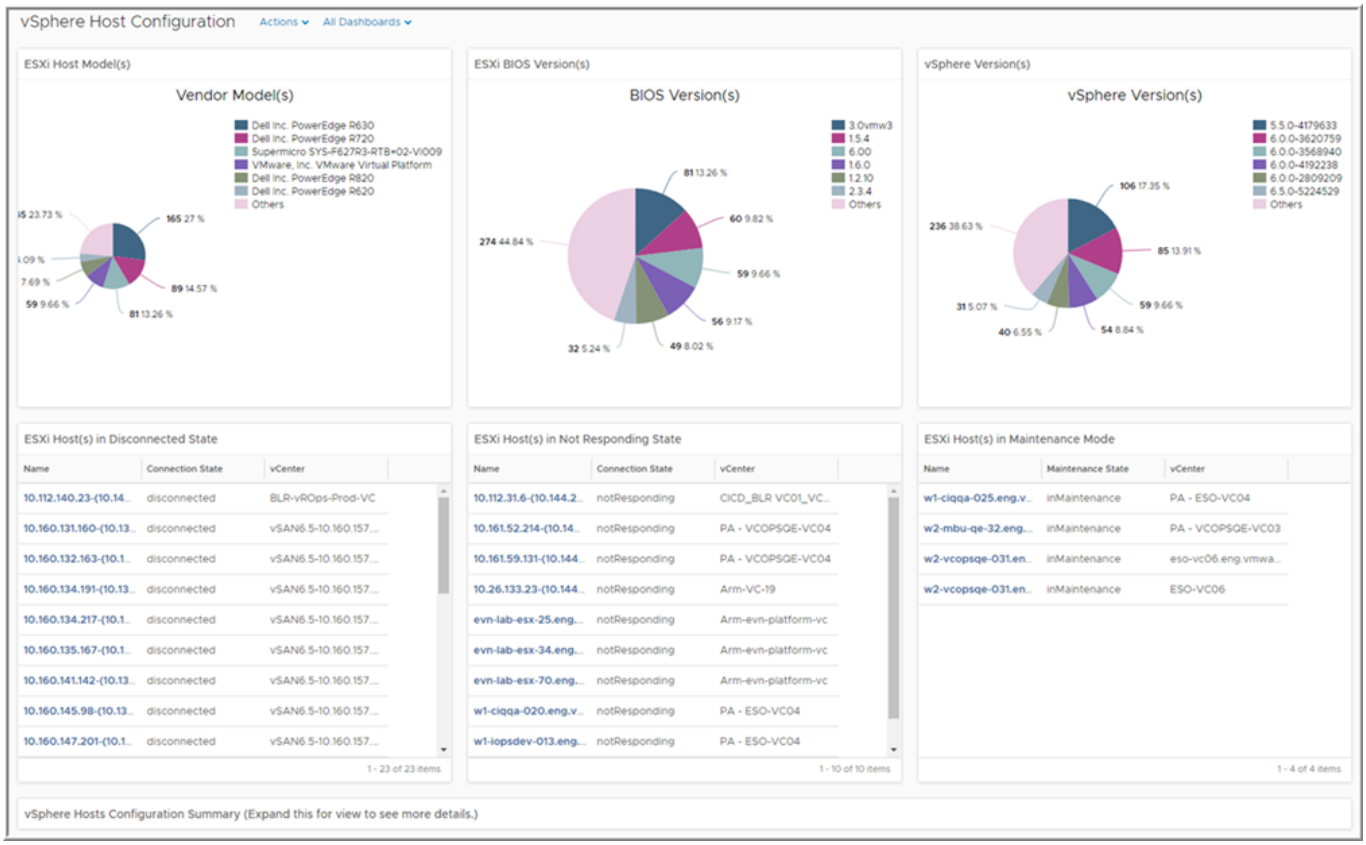


The dashboard also measures the ESXi hosts against the vSphere best practices and indicates deviations that can impact the performance or availability of your virtual infrastructure. Although you can view this type of data in other dashboards, in this dashboard you can export the ESXi configuration view and share it with other administrators.

### SDDC Health and Configuration

This category caters to Senior Executives and virtual infrastructure administrator by providing a summary of health and configuration of the SDDC Management Stack.

Health & Configuration dashboard is available under the SDDC Health and Configuration category.

### Health and Configuration Dashboard

This dashboard provides health of all the SDDC components. The AggregatorCloud Federation Analytics adapter monitors and displays the health of vCenter, VMware Aria OperationsVMware Cloud Foundation Operations health, VMware Aria Automation and other stack components.

> **NOTE**
> Verify to install and configure the SDDC health adapter installed on the child VMware Aria OperationsVMware Cloud Foundation Operations instances to display the health and configuration information in this dashboard.

**Figure 18: Health and Configuration Dashboard**


This dashboard provides information of summary of the number of SDDC components, applications, deployments and its health status. The **SDDC instances and Issue count** widget in the dashboard displays detailed information of which object requires an immediate attention and from which instance the issue has occurred. When you click on the particular VMware Aria OperationsVMware Cloud Foundation Operations child instance, you can drill down to analyze the root cause of an issue.

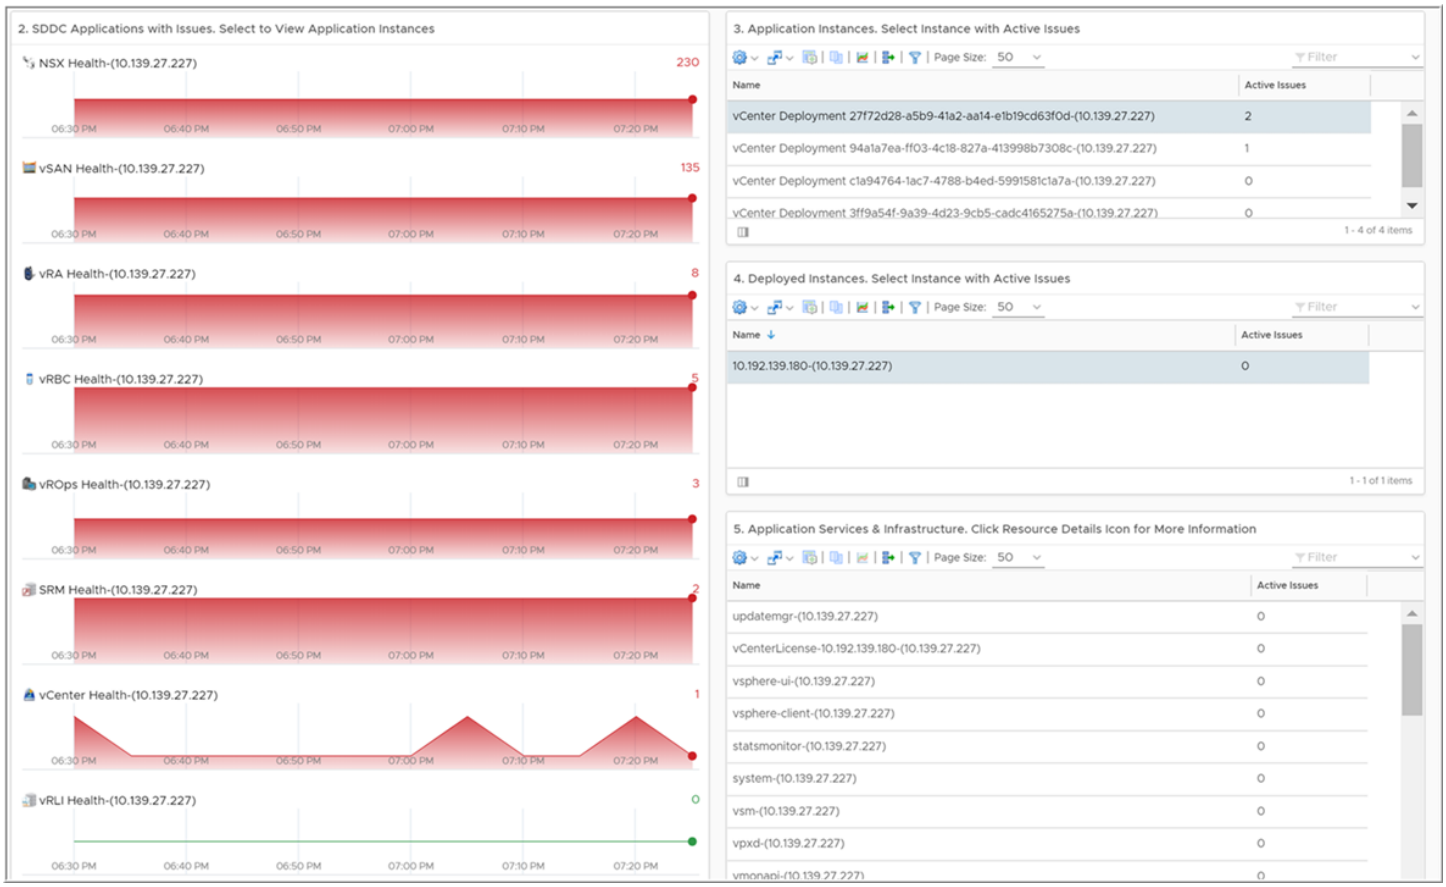


## Dashboards in VMware Infrastructure Health

Use the VMware Infrastructure Health dashboards to view data related to its certificates, historic trend, health, VM, and versions of all products under the VCF management domain.

**Table 196: VMware Infrastructure Health Dashboards**

| Dashboard | Description |
|---|---|
| VMware Infrastructure Health Certificate Overview Dashboard | The VMware Infrastructure Health dashboard provides certificate information related to vCenter Server, NSX Server, NSX Manager Node, ESXi, SDDC Product, and Adapter certificate alerts. |
| VMware Infrastructure Health Historic Trend | Use VMware Infrastructure Health health trend dashboard to view the health trend for each component in the cloud management applications.<br><br>The dashboard provides an overall trend of the health components of the cloud management applications, that are monitored for the last seven days. |
| VMware Infrastructure Health Health Overview | You can use VMware Infrastructure Health dashboard to view and analyze the application-specific problems in the accounts in the VMware Cloud Management Plane. This Dashboard provides the health information for each of the |

*Table continued on next page*

*Continued from previous page*

| Dashboard | Description |
|---|---|
| | accounts. You can select the component in the cloud management applications from the widgets available in the dashboard. The widgets help in rendering the infrastructure health of that component with its service health and associated configuration alerts, if any. |
| VMware Infrastructure Health VM Optimizer | The VMware Infrastructure Management VM Optimizer dashboard helps you to know how you can reclaim capacity from oversized, idle, powered off VMs and optimize these VMs for performance. |
| VMware Infra VCF BOM Version | The VMware Infra VCF BOM Version dashboard displays all the versions of all products under the VCF management domain. |
| VCF HRM Host Systems | The VCF HRM Host System dashboard displays the health, reports health issues, and allows you to manage the health of the VCF host systems. |
| VCF HRM | The VCF HRM dashboard displays the health, reports health issues, and allows you to manage the health of each application in each VCF management domain. |

## vSphere Dashboards

The vSphere dashboard provides an overview of ESXi hosts availability, vCenter appliance availability, and vSphere daily check of the vSphere clusters in your VMware Cloud Foundation Operations environment. The following vSphere dashboards are added to the predefined VMware Aria Operations dashboards:

- ESXi Host Availability
- vCenter Appliance Availability
- vSphere Daily Check

### ESXi Host Availability

ESXi Host Availability dashboard displays the list of unresponsive ESXi hosts, list of hosts in maintainance mode, and list of powered on or unknown power state of the ESXi hosts.

Hosts in maintainance mode and power state is an user initiated action. A host becomes becomes unresponsive because of an external factor that vCenter Server is unaware of. This can happen due to several reasons such as a network connectivity issue between the host and vCenter Server.

On the ESXi Hosts Availability dashboard, the Unresponsive ESXi widget displays the **Power State** as **Unknown** and **Maintainance State** as **notIn Maintainance**.

**How to Use the Dashboard**

- In the Unresponsive ESXi widget, click any ESXi host from the list.
  - The **Parent Cluster of selected ESXi** widget is automatically displayed once you click any host. This widget displays the cluster properties of the selected ESXi hosts such as HA enabled, DRS enabled.
  - The **ESXi Hosts in the selected cluster** widget is automatically displayed once you click any cluster property tab. This widget displays the number of unresponsive and running hosts. You can identify which hosts has issues and try resolving the issues by accessing the hosts.
- Click any host in the **ESXi on Maintenance Mode** or **Power on or Unknown Power State ESXi** widget to view the details in the **Parent Cluster of selected ESXi** widget.

- Click any cluster property tab in the **Parent Cluster of selected ESXi** widget to view the details in the **ESXi Hosts in the selected cluster**.

### vCenter Appliance Availability

vCenter Appliance Availability dashboard displays the list of vCenter instances with reachability and availability issues.

If you have hundreds of vCenter instances and need to check their availability and the status of their services, it can be tedious to log in to each vCenter individually to gather this information. This dashboard helps you identify reachability issues and services that are encountering problems.

#### How to Use the Dashboard

The **vCenter Appliance** widget displays the instance name, total ESXi hosts, VMs, reachability of each vCenter. To check if there are any issues in the vCenter services, click a vCenter instance in the vCenter Appliance widget. The **Services in the selected vCenter** widget is automatically displayed once you click a vCenter instance. This widget displays the health status of the vCenter services. You can sort the services based on health status according to the following status: Started, Stopped, or Unknown.

Clicking an instance shows you the summary of the instance with metrics, logs, alerts, and other details.

### vSphere Daily Check Dashboard

Use the vSphere daily check dashboard to monitor the live operations of the vSphere clusters in your VMware Aria Operations environment.

Use the vSphere Daily Check dashboard to review the health and performance of the vSphere clusters in your VMware Aria Operations environment. The dashboard displays the critical and important alerts that need immediate attention. You can use this dashboard to prevent or minimize alerts for the next 12 to 24 hours. This dashboard focuses on the daily activities of the environment and provides updates based on changes or alerts in the clusters. The dashboard reflects the following changes in your environment:

- Configuration changes
- Consumption changes: Any unexpected increase or sudden decrease in consumption.
- Supply changes. Any unexpected decrease in supply.
- Dynamic changes: Any VM state, VM location, or VM inventory changes.

Use the dashboard to focus on issues that need immediate attention and resolution for proper functioning of your environment. You can adjust the thresholds and parameters to suit your need and add metrics like error logs to customize the dashboard.

This dashboard complements the alerts triggered in your environment and provides insights into the cause of the alerts and also shows the overall picture. This helps troubleshoot the specific issue without disrupting the whole environment. As part of the daily checks, the following areas are covered:

- Availability: Daily checks cover the availability of powered on VMs and ESXi hosts. The numbers reflected should be within the expected range. For example, if a new cluster is added or if several new VMs have been provisioned the day before, the dashboard should reflect the revised available of the VMs and ESXi hosts.
- Performance: Daily checks cover the average performance of consumers (VMs) and providers (compute, network, storage). For example, the dashboard monitors the total CPU, memory, disk, and network utilization as any sudden increase or decrease in performance can cause issues.
- Compliance: Daily checks cover the compliance status and focuses on the present status. The present status covers the past 12 to 24 hours.
- Configuration: Daily checks cover the changes in settings and make sure that the settings match the authorized changes executed during the change window.

# Widgets in VMware Aria OperationsVMware Cloud Foundation Operations

Configuring Widgets
Widgets are the panes on your dashboards. You add widgets to a dashboard to create a dashboard. Widgets display information about attributes, resources, applications, or the overall processes in your environment.

You can configure widgets to reflect your specific needs. The available configuration options vary depending on the widget type. You must configure some of the widgets before they display any data. Many widgets can provide or accept data from one or more widgets. You can use this feature to set the data from one widget as filter and display related information on a single dashboard.

## Widget Interactions

Widget interactions are the configured relationships between widgets in a dashboard where one widget provides information to a receiving widget. When you are using a widget in the dashboard, you select data on one widget to limit the data that appears in another widget, allowing you to focus on a smaller subset data.

### How Interactions Work

If you configured interactions between widget at the dashboard level, you can then select one or more objects in the providing widget to filter the data that appears in the receiving widget, allowing you to focus on data related to an object.

To use the interaction option between the widgets in a dashboard, you configure interactions at the dashboard level. If you do not configure any interactions, the data that appears in the widgets is based on how the widget is configured.

When you configure widget interaction, you specify the providing widget for the receiving widget. For some widgets, you can define two providing widgets, each of which can be used to filter data in the receiving widget.

For example, if you configured the Object List widget to be a provider widget for the Top-N widget, you can select one or more objects in the Object List widget and the Top-N displays data only for the selected objects.

For some widgets, you can define more than one providing widget. For example, you can configure the Metric Chart widget to receive data from a metrics provider widget and an objects providing widget. In such case, the Metric Chart widget shows data for any object that you select in the two provider widgets.

## Configuration Files

You can create configuration files to upload SVG content and define topological hierarchies in dashboards through supported widgets. You can also create configuration files that displays adapter kind, resource kind, and the associated metrics that can be displayed in dashboards using supported widgets.

### Configuration Files for Widget Metric Configuration

You can create an XML file that displays the adapter type, resource kind, and the associated metrics that can be displayed in a widget and dashboard.

### How Configuration Files Work

From the **Configuration Files** page, you can create a XML configuration file with the adapter type, object type, and metric details to be displayed in a dashboard. The final output is an XML file. The supported widgets are Metric Chart, Property List, Rolling View Chart, Scoreboard, Sparkline Chart, and Topology Graph. To use the metric configuration, you must set the widget Self Provider to **Off** and create a widget interaction with a provider widget.

**Where You Find Widget Metric Configuration File**

To manage the configuration files, from the left menu, click **Operations › Configurations**. From the right panel, under **Configuration Files**, click the **Widget Metric Configuration** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

**Table 197: New Configuration File Page Options**

| Option | Description |
|---|---|
| Name | Enter a name for the configuration file. |
| Description | Enter a description for the configuration file. |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | You can define the adapter type, object type, and metrics. From the **Adapter Kind** drop down above the text box, select the adapter type. From the **Resource Kind** drop down above the text box, select the object type. From the **Metric** drop down above the text box, select a metric.<br><br>Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the XML that was created. |

**Table 198: Widget Metric Configuration Toolbar and Data Grid Options**

| Option | Description |
|---|---|
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save** |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them. |
| Import | Select this option if you want to import configuration files. To import:<br>• Click the **Import** option from the horizontal ellipsis.<br>• Click **Browse** and select the file to import. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Select if you want to Overwrite or Skip the file in case of a conflict.<br>• Click **Import** to import the configuration file, and click **Done**. |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| In Use | Indicates whether the configuration file is in use in a dashboard and supported widget.<br>**Yes**: Indicates that the configuration file is used in a dashboard. Hover and click on the green tick mark to view the name of the dashboard and the widget that uses this configuration file.<br><br>**No**: Indicates that the configuration file is not used in a dashboard. |
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, `admin` or `maintenanceAdmin`. |

### Configuration Files for Text Widget Content

You can upload an SVG in Text widgets that are utilized in different dashboards.

#### How Configuration Files Work

From the **Configuration Files** page, you can upload an SVG file to be displayed in a Text widget.

#### Where You Find Text Widget Configuration File

To manage the configuration files, from the left menu, click **Operations › Configurations**. From the right panel, under **Configuration Files**, click the **Text Widget Content** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

**Table 199: New Configuration File Page Options**

| Option | Description |
|---|---|
| Name | Enter a name for the configuration file. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Description | Enter a description for the configuration file. |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | From your system, using an XML editor, open the SVG file to be uploaded. Copy the text from the SVG file into the text box. Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the SVG that was uploaded.<br>**NOTE**<br>Ensure that the text in the text box starts with `<svg`. |

**Table 200: Text Widget Toolbar and Data Grid Options**

| Option | Description |
|---|---|
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save**. |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them. |
| Import | Select this option if you want to import configuration files. To import:<br>• Click the **Import** option from the horizontal ellipsis.<br>• Click **Browse** and select the file to import.<br>• Select if you want to Overwrite or Skip the file in case of a conflict.<br>• Click **Import** to import the configuration file, and click **Done**. |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| In Use | Indicates whether the configuration file is in use in a Text widget and dashboard.<br>**Yes**: Indicates that the configuration file is used in a **Text** widget. Hover and click on the green tick mark to view the name of the dashboard and the Text widget that uses this configuration file.<br><br>**No**: Indicates that the configuration file is not used in a dashboard. |
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, `admin` or `maintenanceAdmin`. |

### Management Packs Configuration

You can create an XML file for specific adapter kinds with additional threshold details such as configuration limits, name, type, and value. The configuration file can be used in a widget and dashboard.

#### How Configuration Files Work

From the **Management Packs Configuration** page, you can create a XML file with adapter specific details with additional thresholds to be displayed in dashboards and widgets. The final output is an XML file. The final output is an XML file. The supported widgets are Metric Chart, Property List, Rolling View Chart, Scoreboard, Sparkline Chart, and Topology Graph.

#### Where You Find Resource Kind Metric Configuration File

To manage the configuration files, from the left menu, click **Operations › Configurations**. From the right panel, under **Configuration Files**, click the **Management Packs Configuration** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

**Table 201: New Configuration File Page Options**

| Option | Description |
|---|---|
| Name | Enter a name for the configuration file. |
| Description | Enter a description for the configuration file. |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | You can define the adapter kind and thresholds such as configuration limits, name, type, and value.<br>Select the tags above the text box to add to the configuration file.<br><br>Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the XML that was created. |

**Table 202: Solutions Configuration Toolbar and Data Grid Options**

| Option | Description |
|---|---|
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save** |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them. |
| Import | Select this option if you want to import configuration files.<br>To import:<br>• Click the **Import** option from the horizontal ellipsis.<br>• Click **Browse** and select the file to import.<br>• Select if you want to Overwrite or Skip the file in case of a conflict.<br>• Click **Import** to import the configuration file, and click **Done**. |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, `admin` or `maintenanceAdmin`. |

### Configuration Files for the Topology Widget

You can create an XML file with topology hierarchies to be displayed in a Topology widget.

#### How Configuration Files Work

From the **Configuration Files** page, you can create a XML file with topology hierarchies to be displayed in a Topology widget. The final output is an XML file.

#### Where You Find Resource Kind Metric Configuration File

To manage the configuration files, from the left menu, click **Operations › Configurations**. From the right panel, under **Configuration Files**, click the **Topology Widget Configuration** tile. To add a configuration file, click the **Add** button. You see the **Create Configuration File** page.

**Table 203: New Configuration File Page Options**

| Option | Description |
|---|---|
| Name | Enter a name for the configuration file. |
| Description | Enter a description for the configuration file. |
| Containing Folder | Select the folder under which you want to store the new configuration file. If you have not created a folder, you can select the **User Defined** folder if it appears. The **System Defined** folder will not appear for selection. |
| Text box | You can define the topology hierarchies.<br>Select the tags above the text box to add to the configuration file.<br><br>Use the orange **Format XML** button to format the XML content. Click **Save**. You see the configuration file under the selected folder. You can preview the XML that was created. |

**Table 204: Topology Widget Configuration Toolbar and Data Grid Options**

| Option | Description |
|---|---|
| Filter | Click the filter icon to filter the configuration files using the following criteria: Name, Description, Last Modified, and Modified By. |
| **Options from the horizontal ellipsis** | |
| Edit | Select a configuration file to edit the contents of the configuration file. |
| Delete | Select one or more configuration files if you want to delete them. You cannot delete configuration files under the **System Defined** folder. |
| Clone | Select a configuration file to clone. You can edit the configuration file options in the **New Configuration File** page and click **Save** |
| Move | Select one or more configuration files if you want to move them to another folder. You can also move configuration files using the drag and drop option. You cannot move a configuration file to or from the **System Defined** folder. |
| Export | Select one or more configuration files if you want to download them |
| Import | Select this option if you want to import configuration files. To import:<br>• Click the **Import** option from the horizontal ellipsis.<br>• Click **Browse** and select the file to import.<br>• Select if you want to Overwrite or Skip the file in case of a conflict.<br>• Click **Import** to import the configuration file, and click **Done**. |
| Create New Folder | Enter a name for the folder and click **OK**. You can create a new folder for meaningful grouping. |
| Rename Folder | Select a folder created by a user and rename the folder. You cannot rename the **System Defined** and **User Defined** folders. |
| **Data Grid Options** | |
| Name | Lists the configuration files that have been created. The **System Defined** folder contains configuration files that are predefined. The **User Defined** folder contains orphan configuration files that were created by users but not assigned to a folder. |
| In Use | Indicates whether the configuration file is in use in a dashboard and topology widget.<br>**Yes**: Indicates that the configuration file is used in a dashboard. Hover and click on the green tick mark to view the name of the dashboard and the widget that uses this configuration file.<br><br>**No**: Indicates that the configuration file is not used in a dashboard. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Description | Displays the description entered while creating/editing the configuration file. |
| Last Modified | Displays when the configuration file was last updated. |
| Modified By | Displays who last modified the configuration file. For example, `admin` or `maintenanceAdmin`. |

## Widget Definitions List

A widget is a pane on a dashboard that contains information about configured attributes, resources, applications, or the overall processes in your environment. Widgets can provide a holistic, end-to-end view of the health of all the objects and applications in your enterprise. If your user account has the necessary access rights, you can add and remove widgets from your dashboards.

**Table 205: Summary of Widgets**

| Widget Name | Description |
|---|---|
| Alert List | Shows a list of alerts for the objects that the widget is configured to monitor. If no objects are configure, the list displays all alerts in your environment. |
| Alert Volume | Shows a trend report for the last seven days of alerts generated for the objects it is configured to monitor. |
| Anomalies | Shows a chart of the anomalies count for the past 6 hours. |
| Anomaly Breakdown | Shows the likely root causes for symptoms for a selected resource. |
| Capacity Remaining | Shows a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource. |
| Container Details | Shows the health and alert counts for each tier in a single selected container. |
| Container Overview | Shows the overall health and the health of each tier for one or more containers. |
| Current Policy | Shows the highest priority policy applied to a custom group. |
| Data Collection Results | Shows a list of all supported actions specific for a selected object. |
| DRS Cluster Settings | Shows the workload of the available clusters and the associated hosts. |
| Efficiency | Shows the status of the efficiency-related alerts for the objects that it is configured to monitor. Efficiency is based on generated efficiency alerts in your environment. |
| Environment | Lists the number of resources by object or groups them by object type. |
| Environment Overview | Shows the performance status of objects in your virtual environment and their relationships. You can click an object to highlight its related objects and double-click an object to view its Resource Detail page. |
| Environment Status | Shows statistics for the overall monitored environment. |
| Faults | Shows a list of availability and configuration issues for a selected resource. |
| Forensics | Shows how often a metric had a particular value, as a percentage of all values, within a given time period. It can also compare percentages for two time periods. |
| Geo | Shows where your objects are located on a world map, if your configuration assigns values to the Geo Location object tag. |
| Health | Shows the status of the health-related alerts for the objects that it is configured to monitor. Health is based on generated health alerts in your environment. |
| Health Chart | Shows health information for selected resources, or all resources that have a selected tag. |
| Heat Map | Shows a heat map with the performance information for a selected resource. |

*Table continued on next page*

*Continued from previous page*

| Widget Name | Description |
|---|---|
| Mashup Chart | Brings together disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs). This widget is typically used for a container. |
| Metric Chart | Shows a chart with the workload of the object over time based on the selected metrics. |
| Metric Picker | Shows a list of available metrics for a selected resource. It works with any widget that can provide resource ID. |
| Object List | Shows a list of all defined resources. |
| Object Relationship | Shows the hierarchy tree for the selected object. |
| Object Relationship (Advanced) | Shows the hierarchy tree for the selected objects. It provides advanced configuration options. |
| Property List | Shows the properties and their values of an object that you select. |
| Recommended Actions | Displays recommendations to solve problems in your vCenter instances. With recommendations, you can run actions on your data centers, clusters, hosts, and virtual machines. |
| Risk | Shows the status of the risk-related alerts for the objects that it is configured to monitor. Risk is based on generated risk alerts in your environment. |
| Rolling View Chart | Cycles through selected metrics at an interval that you define and shows one metric graph at a time. Miniature graphs, which you can expand, appear for all selected metrics at the bottom of the widget. |
| Scoreboard | Shows values for selected metrics, which are typically KPIs, with color coding for defined value ranges. |
| Scoreboard Health | Shows color-coded health, risk, and efficiency scores for selected resources. |
| Sparkline Chart | Shows graphs that contain metrics for an object . If all the metrics in the Sparkline Chart widget are for an object that another widget provides, the object name appears at the top right of the widget. |
| Tag Picker | Lists all defined resource tags. |
| Text Display | Reads text from a Web page or text file and shows the text in the user interface. |
| Time Remaining | Shows a chart of the Time Remaining values for a specific resource over the past 7 days. |
| Top Alerts | Lists the alerts most likely to negatively affect your environment based on the configured alert type and objects. |
| Top-N | Shows the top or bottom N number metrics or resources in various categories, such as the five applications that have the best or worst health. |
| Topology Graph | Shows multiple levels of resources between nodes. |
| View | Shows a defined view depending on the configured resource. |
| Weather Map | Uses changing colors to show the behavior of a selected metric over time for multiple resources. |
| Workload | Shows workload information for a selected resource. |
| Workload Pattern | Shows a historical view of the hourly workload pattern of an object. |

For more information about the widgets, see the VMware Aria OperationsVMware Cloud Foundation Operations help.

**Alert List Widget**

The Alert List widget is a list of alerts for the objects it is configured to monitor. You can create one or more alert lists in VMware Aria OperationsVMware Cloud Foundation Operations for objects that you add to your custom dashboards. The widget provides you with a customized list of alerts on objects in your environment.

## How the Alert List Widget and Configuration Options Work

You can add the Alert List widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. You edit an Alert List widget after you add it to a dashboard. The changes you make to the options create a custom alert list to meet the needs of the dashboard users.

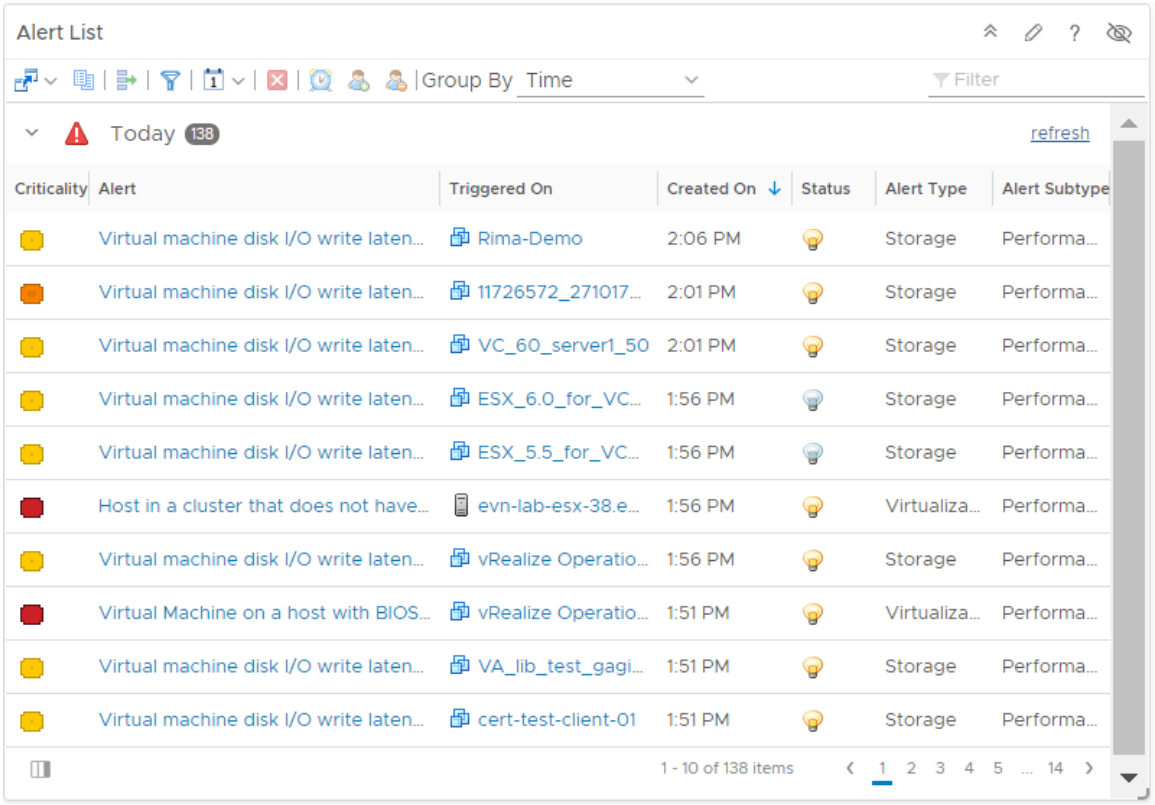


## Where You Find the Alert List Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Alert List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Dashboard Navigation | Actions you can run on the selected alert. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | For example, you use the option to open a vCenter, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems. |
| Reset Interaction | Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.<br><br>Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards. |
| Perform Multi-Select Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br><br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |
| Display Filtering Criteria | Displays the object information on which this widget is based. |
| Select Date Range | Limits the alerts that appear in the list to the selected date range. |
| Cancel Alert | Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.<br><br>You cancel alerts when you do not need to address them. Canceling the alert does not cancel the underlying condition that generated the alert. Canceling alerts is effective if the alert is generated by triggered fault and event symptoms because these symptoms are triggered again only when subsequent faults or events occur on the monitored objects. If the alert is generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again. |
| Suspend | Suspend an alert for a specified number of minutes.<br><br>You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object.<br><br>The user who suspends the alert becomes the assigned owner. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>You can cancel or retrigger the alert, if it is still active when its suspension period has ended, by rerunning the automated actions connected to the alert. In this case, you can suppress cancelation and update on all instances of an alert on an object. To activate this option, open the property file `/usr/lib/vmware-vcops/user/conf/analytics/advanced.properties` and add `retriggerExpiredSuspendedActiveAlerts = true` to the property file, and restart the VMware Aria Operations analytics service or the VMware Aria Operations cluster. |
| Take Ownership | As the current user, you make yourself the owner of the alert.<br><br>You can only take ownership of an alert, you cannot assign ownership. |
| Release Ownership | Alert is released from all ownership. |
| Group By | Group alerts by the options in the drop-down menu. |
| Filter | Locate data in the widget. |

**Table 206: Group By Options**

| Option | Description |
|---|---|
| None | Alerts are not sorted into specific groupings. |
| Time | Group alerts by time triggered. The default. |
| Criticality | Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the Alert List Widget Data Grid table. |
| Definition | Group alerts by definition, that is, group like alerts together. |
| Object Type | Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together. |

**Alert List Widget Data Grid Options**

The data grid provides information on which you can sort and search.

Expand the grouped alerts to view the data grid.

| Option | Description |
|---|---|
| Criticality | Criticality is the level of importance of the alert in your environment. The alert criticality appears in a tooltip when you hover the mouse over the criticality icon. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**. |
| Alert | Description of the alert. |
| Triggered On | Name of the object for which the alert was generated. |
| Created On | Date and time when the alert was generated. |
| Status | Current state of the alert. |
| Alert Type | Alert type is assigned when you create the alert definition. It helps you categorize and route the alert to the appropriate domain administrator for resolution.<br><br>The possible values include:<br>• Application<br>• Virtualization/Hypervisor<br>• Hardware (OSI)<br>• Storage<br>• Network |
| Alert Sub-Type | Alert subtype is assigned when you create the alert definition. It helps you categorize and route the alert to the appropriate domain administrator for resolution.<br><br>The possible values include:<br>• Availability<br>• Performance<br>• Capacity<br>• Compliance<br>• Configuration |
| Importance | Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm. |

**Alert List Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br><br>If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsectio n, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |
| Alert Related | A group of filters limits the alerts that appear in this alert list to those that meet the selected criteria.<br><br>If the objects on which the alerts are based have an input transformation applied, you define filters for the alerts based on the transformed objects.<br><br>You can configure the following filters:<br>• Status. Select one or more alert states to include in the list.<br>• Criticality. Select one or more levels of criticality. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Control State. Select one or more control states to include in the list.<br>• Impact. Select one or more alert badges to include in the list.<br>• Actions.<br>• Alert Type. Select the subtype in the type list. This value was assigned when you configured the alert definition.<br>• Alert Definition. Drag and drop the alert definitions to the left pane from the Alert Definitions list. Click **OK** to filter by the selected Alert Definitions.<br><br>The selections you make for each filter are displayed as labels for easy visibility. |

### Alert Volume Widget

The Alert Volume widget is a trend report for the last seven days of alerts generated for the objects it is configured to monitor in VMware Aria OperationsVMware Cloud Foundation Operations. You can create one or more alert volume widgets for objects that you add to your dashboards. The alert volume provides you with a customized trend report on objects that helps you identify changes in alert volume, indicating a problem in your environment.

#### How the Alert Volume Widget and Configuration Options Work

You can add the Alert Volume widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

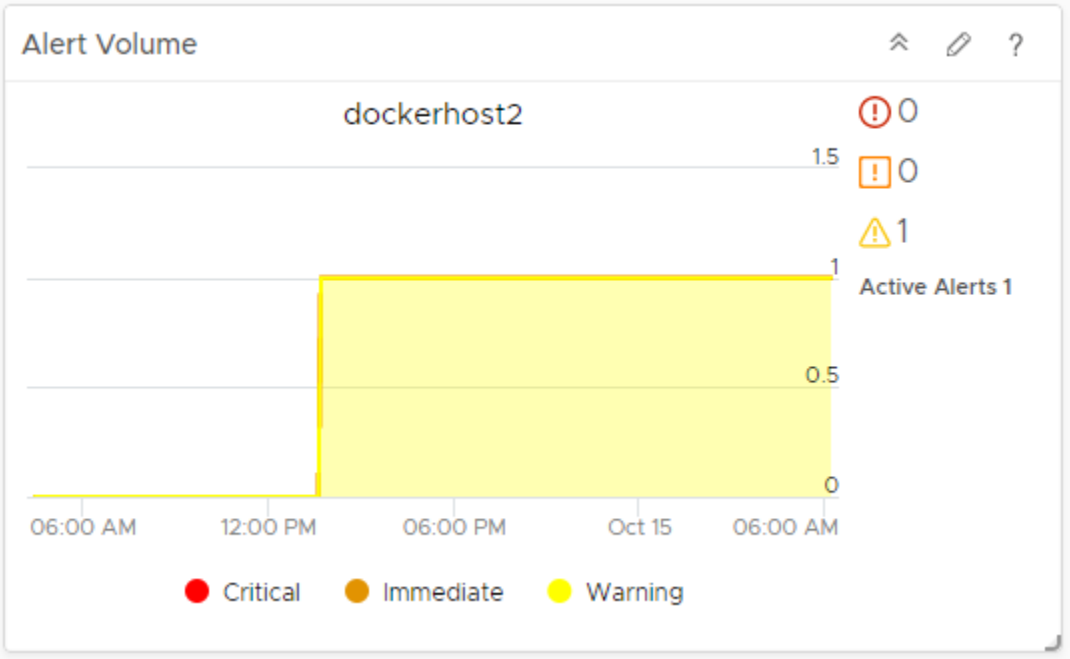

### Where You Find the Alert Volume Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Alert Volume Widget Display Options

The Alert Volume widget displays a trend chart, symptoms by criticality, and active alerts.

| Option | Description |
|---|---|
| Trend chart | Volume of critical, immediate, and warning symptoms for the configured objects. |
| Symptoms by criticality | Number of symptoms for each criticality level. |
| Active Alerts | Number of active alerts. Alerts can have more than one triggering symptom. |

### Alert Volume Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

**Anomalies Widget**

The Anomalies widget displays the anomalies for a resource for the past 6 hours at time intervals you set.

The Anomalies widget shows or hides time periods when the metric violates a threshold that is configured. The widget color indicates the criticality of the violation.

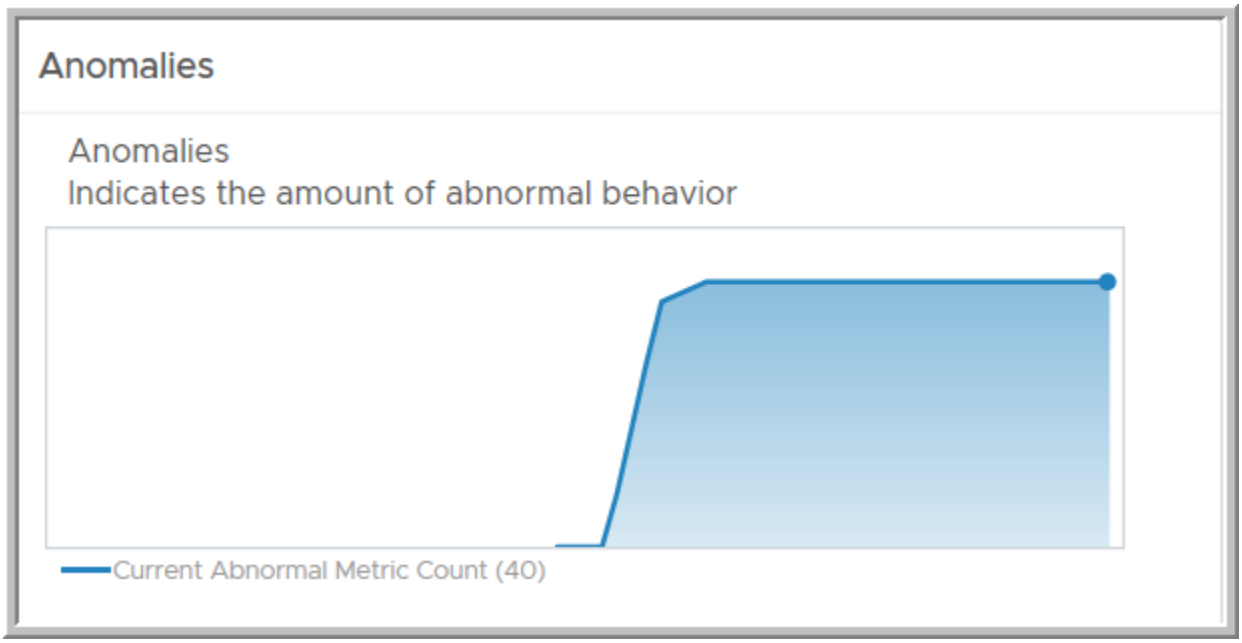


**Where You Find the Anomalies Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a

widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Anomalies Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | • On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

### Anomaly Breakdown Widget

The Anomaly Breakdown widget shows the likely root causes for symptoms for a selected resource.

## How the Anomaly Breakdown Widget and Configuration Options Work

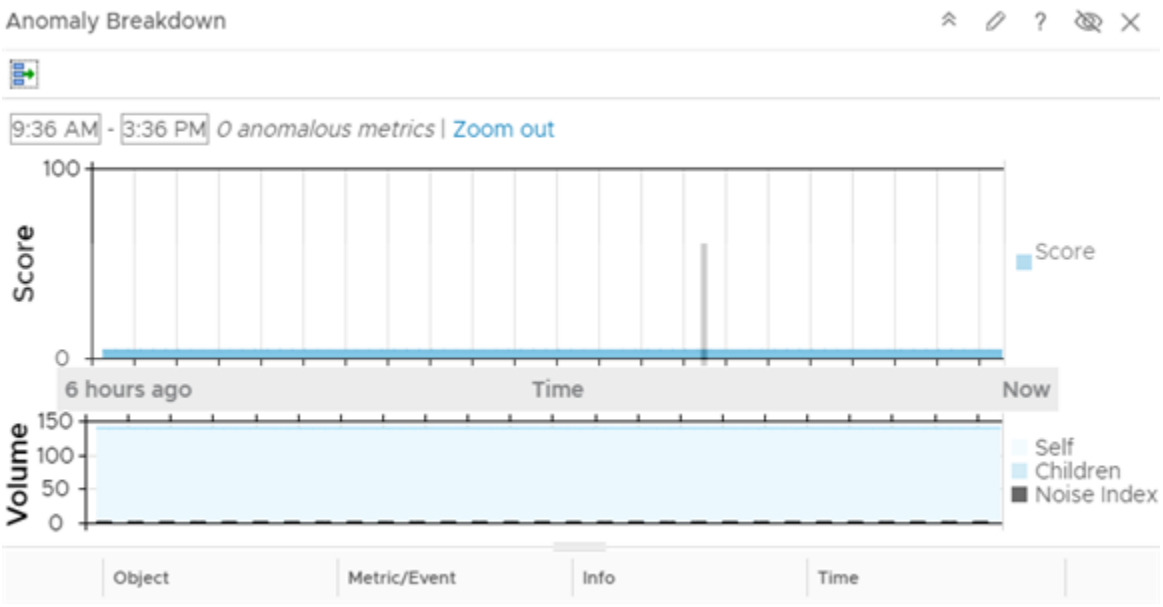


You can add the Anomaly Breakdown widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

### Where You Find the Anomaly Breakdown Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Anomaly Breakdown Widget Display Options

The Anomaly Breakdown widget displays scores, volume, and a list of anomaly metrics.

| Option | Description |
|---|---|
| Score | Anomaly value. |
| Volume | VMware Aria OperationsVMware Cloud Foundation Operations full set metric count for the selected object in the specified time range. |
| Anomaly Metrics List | List of alarms for the selected object in the specified time range. |

### Anomaly Breakdown Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Show Bar Details | If the widget is displaying data for multiple objects, you can select a row and click this button to view the list of alarms for the selected object. |
| Perform Multiple Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br><br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |

**Anomaly Breakdown Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Mode | Display a single object or multiple objects. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Show | Select the number of objects to display in multiple objects mode. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |

### Container Details Widget

The Container Details widget displays graphs that show a summary of child objects, metrics, and alerts of an object in the inventory.

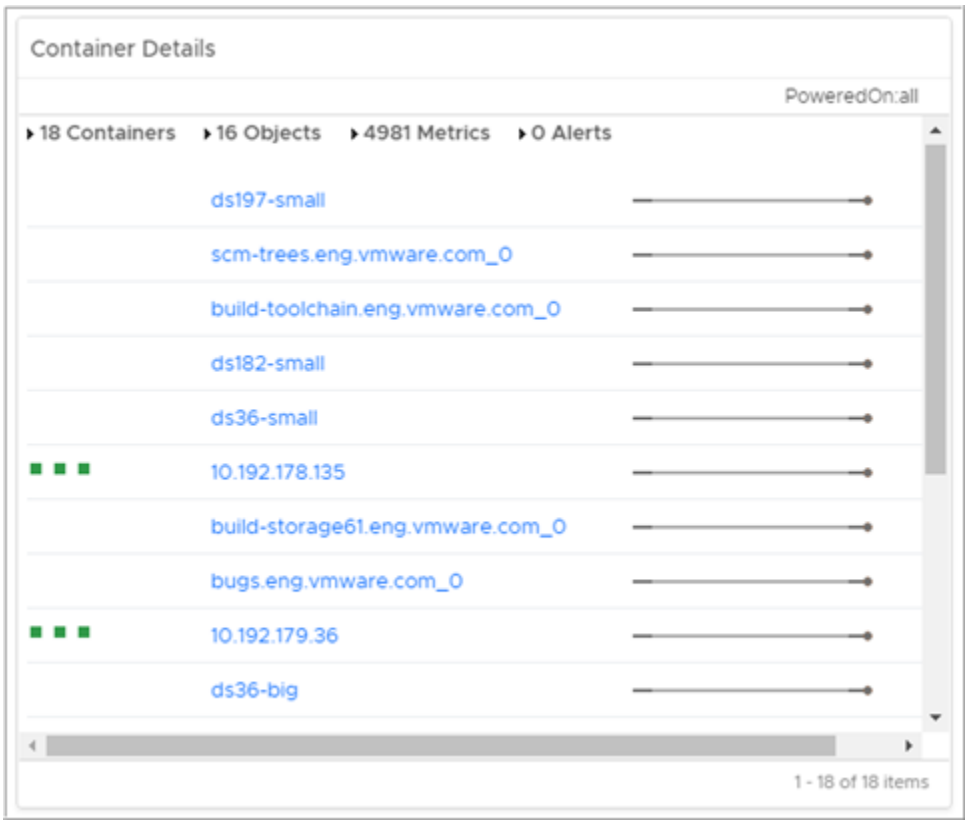

## How the Container Details Widget and Configuration Options Work

The Container Details widget treats objects from the inventory as containers and objects. Containers are objects that contain other objects. The widget lists the containers and shows the number of containers, objects, metrics, and alerts of the observed object. The widget also displays the alerts of each container and an icon links to its child objects. For example, if you select from the inventory a host that contains three objects such as, two virtual machines and one datastore, the Container Details widget displays summary information with three containers, two objects that are the child objects of the two virtual machines, and the number of alerts for the host and the number of metrics for the child objects of the host. The widget also lists each of the three containers, with the number of alerts for each object. Clicking an object in the graph takes you to the object details page. When you point to the icon next to the object, a tool tip shows the name of the related resource and its health. For example, when you point to the icon next to a virtual machine, the tool tip shows a related datastore and its health. Clicking the icon takes you to the object detail page of the related object, which is the datastore following the example.

You edit a container details widget after you add it to a dashboard. You can configure the widget to take information from another widget in the dashboard and to analyze it. When you select **Off** from the Self Provider option and set source and receiver widgets in the **Widget Interactions** menu during editing of the dashboard, the receiver widget shows information about an object that you select from the source widget. For example, you can configure the Container Details widget to display information about an object that you select from the Object Relationship widget in the same dashboard.

## Where You Find the Container Details Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Container Details Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Mode | You can change the size of the graph using the Compact or Large buttons. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

**Capacity Remaining Widget**

The Capacity Remaining widget displays a percentage indicating the remaining computing resources as a percent of the total consumer capacity. It also displays the most constrained resource.

**Where You Find the Capacity Remaining Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Capacity Remaining Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filte r** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

### Container Overview Widget

The Container Overview widget gives a graphical presentation of the health, risk, and efficiency of an object or list of objects in the environment.

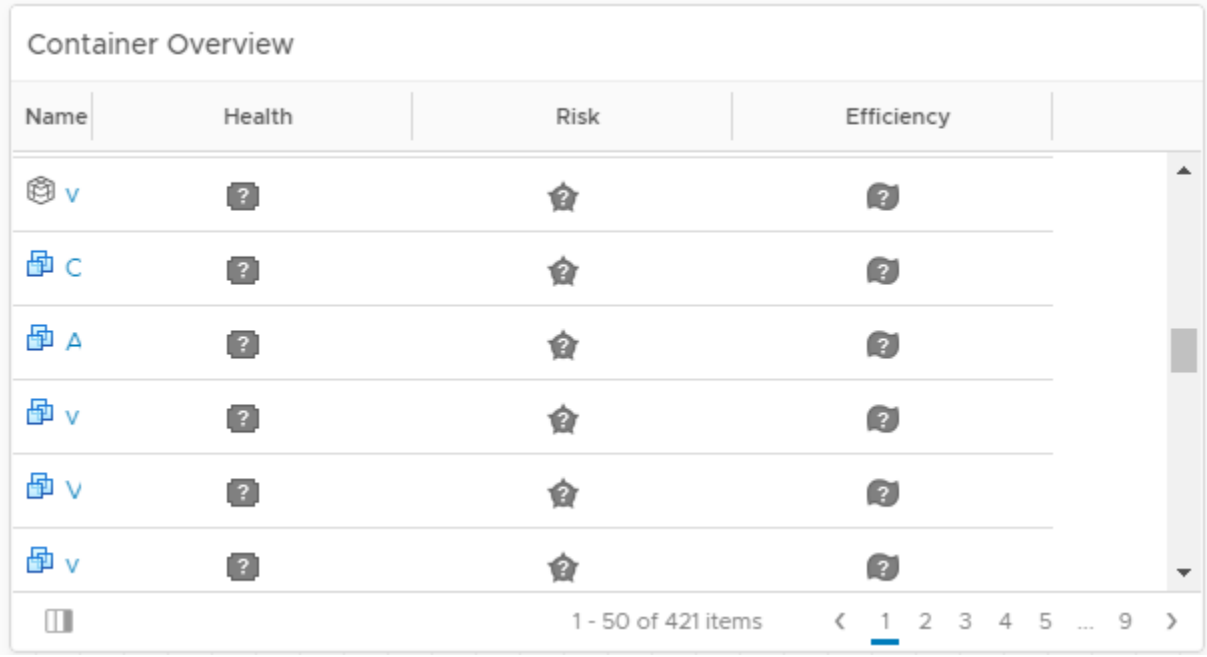


#### How the Container Overview Widget and Configuration Options Work

The Container Overview widget displays the current status, the status for a previous time period of the health, risk, and the efficiency of an object or list of objects. You can configure the widget to display information for one or more objects that you are interested in when you select the **Object** mode during configuration of the widget. The widget displays

information for all objects from an object type or types when you select the **Object Type** mode during configuration of the widget. You can open the object detailed page of each object in the data grid when you click the object.

You edit a container overview widget after you add it to a dashboard. You can configure the widget to display information about an object or to display information about all objects from an object type by using the **Object** or **Object Type** mode. The configuration options change depending on your selection of mode.

### Where You Find the Container Overview Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Container Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about other widgets or dashboards.

| Option | Description |
|---|---|
| Perform Multi-Select Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br><br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |
| Filter | You can filter the objects in the data grid. |
| Dashboard Navigation | You can explore information from another dashboard.<br><br>**NOTE**<br>This toolbar icon exists when you configure the widget to interact with a widget from another dashboard. Use **Dashboard Navigation** menu during dashboard configuration to configure the widgets to interact.<br><br>When you select an object from an object data grid and click the toolbar icon, it takes you to a related dashboard. For example, you can configure the widget to send information to a Topology Graph widget that is on another dashboard, for example dashboard 1. When you select a VM from the data grid, click **Perform Multi-Select Interaction** , click **Dashboard Navigation** and select **Navigate › dashboard 1**. It takes you to dashboard 1, where you can observe selected VM and objects related to it. |

### Container Overview Widget Data Grid Options

The data grid provides information on which you can sort and search.

| Option | Description |
|---|---|
| Name | Name of the object |
| Health | Shows information about the health parameter.<br><br>Status displays the badge of the current health status of an object. You can check the status in a tool tip when you point to the badge.<br><br>Last 24 Hours displays the statistic of health parameter for last 24 hours. |
| Risk | Shows information about the risk parameter.<br><br>Status displays the badge of the current risk status of an object. You can check the status in a tool tip when you point to the badge.<br><br>Last Week displays the statistics of the health parameter for the last week. |
| Efficiency | Shows information about the efficiency parameter.<br><br>Status displays the badge of the current efficiency status of an object. You can check the status in a tool tip when you point to the badge.<br><br>Last Week displays statistic of the efficiency parameter for the last week. |

### Container Overview Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Mode | Use **Object** to select an object from the environment to observe.<br><br>Use **Object Type** to select the type of the objects to observe. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| Object Type | Select an object type in your environment on which you want to base the widget data.<br>1. Click the **Add Object Type** icon to search for and add an object type.<br><br>When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 2. Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type. |

**Current Policy Widget**

The Current Policy widget displays the active operational policy that is assigned to your object or object group. VMware Aria OperationsVMware Cloud Foundation Operations uses the assigned policy to analyze your objects, control the data that is collected from those objects, generate alerts when problems occur, and display the results in the dashboards.

**How the Current Policy Widget and Configuration Options Work**

You add the Current Policy widget to a dashboard so that you can quickly see which operational policy is applied to an object or object group. To add the widget to a dashboard, you must have access permissions associated with the roles assigned to your user account.

The configuration changes that you make to the widget creates a custom instance of the widget that you use in your dashboard to identify the current policy assigned to an object or object group. When you select an object on the dashboard, the policy applied to the object appears in the Current Policy widget, with an embedded link to the policy details. To display the inherited and local settings for the applied policy, click the link.

**Where You Find the Current Policy Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Current Policy Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.<br><br>For example, to view the policy applied to each object that you select in the Object List widget, select **Off** for Self Provider. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pa ne to select an object based on tag values. |

**Data Collection Results Widget**

The Data Collection Result widget shows a list of all supported actions specific for a selected object. The widget retrieves data specific to a selected object actions and uses the action framework to run data collection actions.

**How the Data Collection Results Widget and Configuration Options Work**

You can add the Data Collection Results widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The Data Collection Results widget is a receiver of a resource or metric ID. It can interact with any resource or metric ID that provides widgets such as Object List and Metric Picker. To use the widget, you must have an environment that contains the following items.

- A vCenter Adapter instance
- A VMware Aria OperationsVMware Cloud Foundation Operations for Horizon View Adapter
- A VMware Aria OperationsVMware Cloud Foundation Operations for Horizon View Connection Server

You edit a Data Collection Result widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

**Where You Find the Data Collection Results Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a

widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Data Collection Results Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Results | Shows all finished and currently running actions for the selected object. |
| Choose Action | Shows a list with all supported actions specific for the selected object. The selected object is a result of widget interactions. |

### Data Collection Results Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget updates only when you open the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Config | Specifies self provider choice and selection of a resource instance. |
| Selected Object | When you select an object, this text box is populated by the object. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Start new data collection on interaction change | Indicates whether to start a new data collection action when the object selection changes in the source widget. |
| Objects | List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget. |
| Defaults | Specifies the default data collection action selected for each object type. |
| Object Types | List of object types in your environment that you can search or sort by column so that you can locate the object type on which you are basing the data that appears in the widget. You can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box. |
| Default Data Collection Action | This panel is populated by the object type that you select in the object types list.<br><br>You can select only one default data collection action for an object type. |

**DRS Cluster Settings Widget**

The DRS Cluster Settings widget displays the workload of the available clusters and the associated hosts. You can change the Distributed Resource Scheduler (DRS) automation rules for each cluster.

**How the DRS Cluster Settings Widget and Configuration Options Work**

You can view CPU workload and memory workload percentages for each of the clusters. You can view CPU workload and memory workload percentages for each host in the cluster by selecting a cluster in the data grid. The details are displayed in the data grid below. You can set the level of DRS automation and the migration threshold by selecting a cluster and clicking **Cluster Actions › Set DRS Automation**.

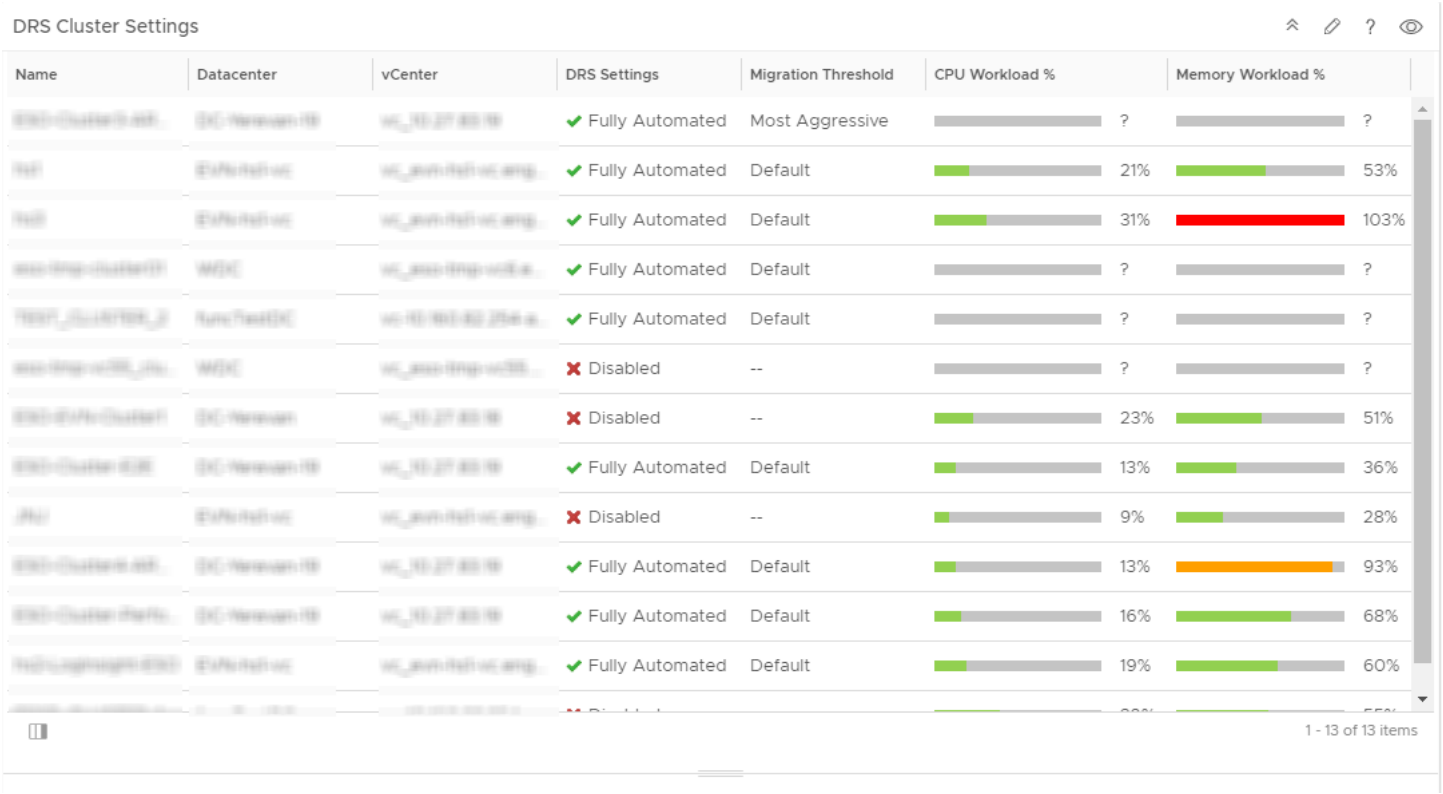


You edit a DRS Cluster Settings widget after you add it to a dashboard. To configure the widget, click the edit icon at the upper-right corner of the widget window. You can add the DRS Cluster Settings widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The DRS Cluster Settings widget appears on the dashboard named vSphere DRS Cluster Settings, which is provided with VMware Aria OperationsVMware Cloud Foundation Operations.

## Where You Find the DRS Cluster Settings Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## DRS Cluster Settings Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Cluster Actions | Limits the list to actions that match the cluster you select. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Show | The drop-down menu displays the parent vCenter instances where the clusters reside. You can also view the data centers under each parent vCenter instance. Select a parent vCenter to view the workload of the available clusters in the data grid.<br><br>The default setting displays the clusters across all vCenters. |
| Filter | Filters the data grid by name, data center, vCenter, DRS settings, and migration threshold. |

**DRS Cluster Settings Widget Data Grid Options**

The data grid provides information on which you can sort and search.

| Option | Description |
|---|---|
| Name | Displays the names of the clusters in the selected parent vCenter instance. |
| Datacenter | Displays the data centers that belong to each cluster. |
| vCenter | Displays the parent vCenter instance where the cluster resides. |
| DRS Settings | Displays the level of DRS automation for the cluster.<br><br>To change the level of DRS automation for the cluster, select **Cluster Actions › Set DRS Automation** from the toolbar. You can change the automation level by selecting an option from the drop-down menu in the Automation Level column. |
| Migration Threshold | Recommendations for the migration level of virtual machines. Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster. |
| CPU Workload % | Displays the percentage of CPU in GHz available on the cluster. |
| Memory Workload % | Displays the percentage of memory in GB available on the cluster. |

**DRS Cluster Settings Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

**Efficiency Widget**

The efficiency widget is the status of the efficiency-related alerts for the objects it is configured to monitor. Efficiency alerts in VMware Aria OperationsVMware Cloud Foundation Operations usually indicate that you can reclaim resources. You can create one or more efficiency widgets for objects that you add to your custom dashboards.

**How the Efficiency Widget and Configuration Options Work**

You can add the efficiency widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning efficiency alerts generated over time, if the monitored object is a group.
- A trend line displays the efficiency status of the monitored object over time if the object does not provide its resources to any other object, or where no other object depends on the monitored object's resources. For example, if the monitored object is a virtual machine or a distributed switch.
- A pie chart displays the reclaimable, stress, and optimal percentages for the virtual machines that are descendants of the monitored object for all other object types. You use the chart to identify objects in your environment from which you can reclaim resources. For example, if the object is a host or datastore.

If the **Badge Mode** is set to **On**, only the badge appears.

Edit an efficiency widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

**Where You Find the Efficiency Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a

widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Efficiency Widget Display Options

The Efficiency widget displays an efficiency badge. The widget also displays an efficiency trend when not in badge mode.

| Option | Description |
|---|---|
| Efficiency Badge | Status of the objects configured for this instance of the widget.<br><br>Click the badge to open the **Alerts** tab for the object that provides data to the widget. |
| Efficiency Trend | Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears. |

### Efficiency Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Badge Mode | Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.<br><br>Select one of the following options:<br>• On. Only the badge appears in the widget.<br>• Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

### Environment Widget

The Environment widget displays the resources which collects data. You can create one or more lists in VMware Aria OperationsVMware Cloud Foundation Operations for the resources that you add to your custom dashboards.

#### How the Environment Widget and Configuration Options Work

The Environment widget lists the number of resources by object or groups them by object type. You can add the Environment widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an Environment widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

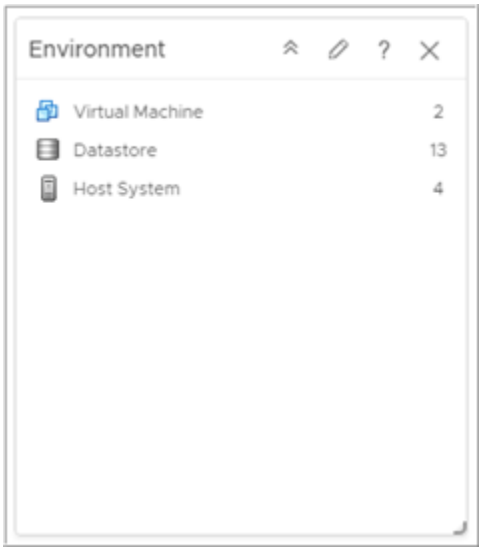

## Where You Find the Environment Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Environment Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

**Environment Overview Widget**

The Environment Overview widget displays the health, risk, and efficiency of resources for a given object from the managed inventory.

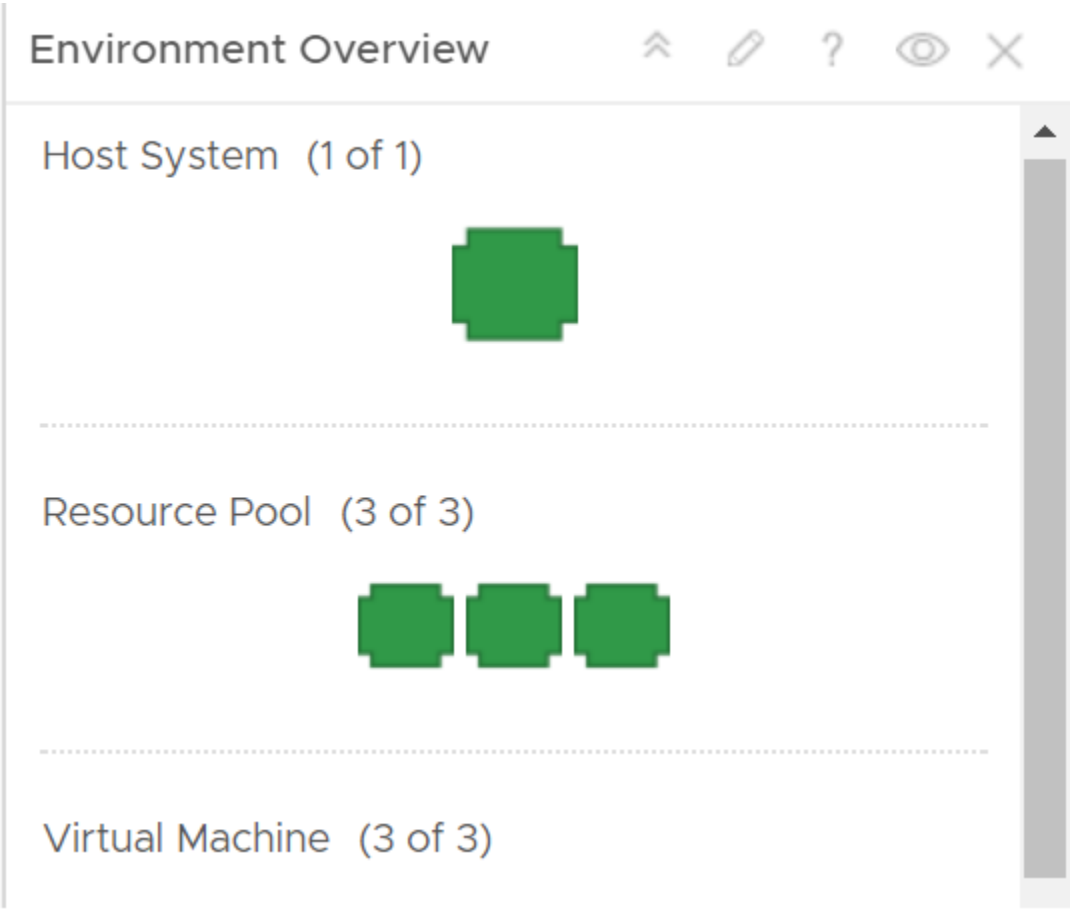


**How the Environment Overview Widget and Configuration Options Work**

You can add the Environment Overview widget to one or more custom dashboards.

The widget displays data for objects from one or several types. The data that the widget displays depends on the object type and category that you selected when you configured the widget.

The objects in the widget are ordered by object type.

The parameters for the health, risk, and efficiency of an object appear in a tool tip when you point to the object.

When you double-click an object on the Environment Overview widget, you can view detailed information for the object.

To use the Environment Overview widget, you must add it to the dashboard and configure the data that appears in the widget. You must select at least one badge and an object. Additionally, you can select an object type.

The Environment Overview widget has basic and advanced configuration options. The basic configuration options are activated by default.

To use all features of the Environment Overview widget, you must change the default configuration of the widget. Log in to the VMware Aria OperationsVMware Cloud Foundation Operations machine and set `skittlesCustomMetricAllowed`

to `true` in the `web.properties` file. The `web.properties` file is located in the `/usr/lib/vmware-vcops/user/conf/web` folder. The change is propagated after you use the **`service vmware-vcops-web restart`** command to restart the UI.

You must use the **Badge** tab to select the badge parameters that the widget shows for each object. You must use the **Config** tab to select an object or object type. To observe a concrete object from the inventory, you can use the **Basic** option. To observe a group of objects or objects from different types, you must use the **Advanced** option.

### Where You Find the Environment Overview Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Environment Overview Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to get more information about badges.

| Option | Description |
|---|---|
| Badge | You can select a Health, Risk, or Efficiency badge for objects that appear in the widget. The tool tip of a badge shows the standard name of the badge. |
| Status | You can filter objects based on their badge status and their state. |
| Sort | You can sort objects by letter or by number. |

### Environment Overview Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| Selected Object | Object that is the basis for the widget data.<br><br>To populate the text box, select **Config › Basic** and select an object from the list. |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Badge | Defines a parameter to observe. You can select or deselect Health, Risk, and Efficiency parameters using check boxes. Default configuration of the widget selects all badges.<br><br>Select at least one badge parameter. |
| Config | **Basic**<br>List of objects in your environment that you can search or sort by column so that you can locate the object on which you are basing the data that appears in the widget. |
| | **Advanced**<br>You can use Object Types to select a type of the objects to observe information about health, risk, and efficiency. Double-click the object type to select it.<br><br>Use the **Adapter Type** drop-down menu to filter the objects types based on an adapter.<br><br>You can use the **Use vSphere Default** button to observe the main vSphere object types.<br><br>To remove an object type from the list, click **Remove Selected** next to **Use vSphere Default**.<br><br>You can use the **Object Type Categories** menu to select a group or groups of object types to observe.<br><br>You can use the Object tree to select an object to filter the displayed objects. For example, to observe a datastore of a VM, double-click **Datastore** from the **Object Types** menu to select it. Click the datastore when it is in the list of object types, and find the VM in the object tree and select it. To return to your previous configuration of the widget, click **Datastore** from the list of object types and click **Deselect All** in the object tree window.<br><br>The metrics tree and badge data grids are available configuration options only if the default configuration of the widget is changed. To use these configuration options, log in to the VMware Aria OperationsVMware Cloud Foundation Operations machine and set |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | `skittlesCustomMetricAllowed` to `true` in the `web.properties` file. The `web.properties` file is located in the `/usr/lib/vmware-vcops/user/conf/web` folder. |

### Environment Status Widget

The Environment Status widget displays the statistics for the overall monitored environment.

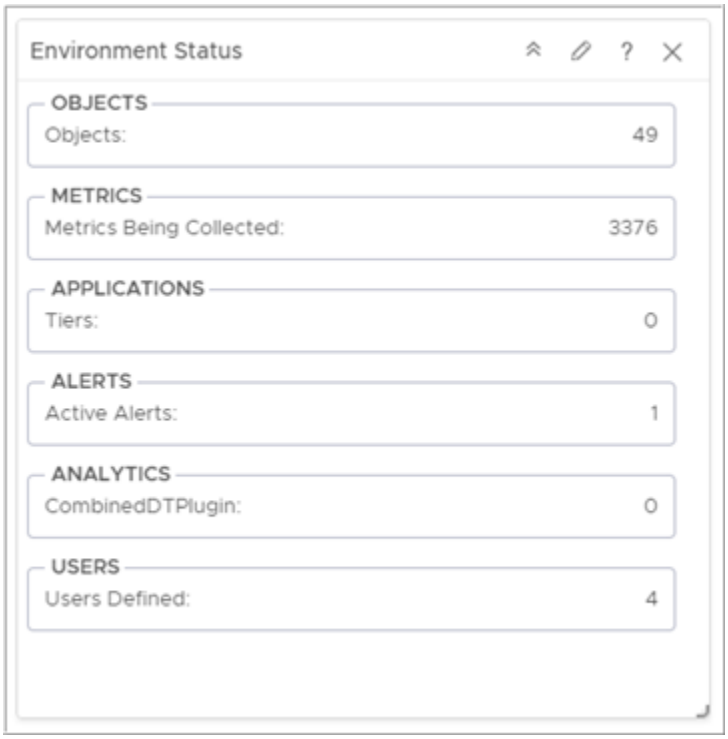


#### How the Environment Status Widget and Configuration Options Work

You customize the output of the widget by choosing a category such as Objects, Metrics, Applications, Alerts, Analytics, and Users. You can filter the data by using the tags tree from **Select which tags to filter** in the configuration window.

You edit an environment status widget after you add it to a dashboard. To configure the widget, click the pencil at the right corner of the widget window. You must select at least one type of information from **OBJECTS**, **METRICS**, **APPLICATIONS**, **ALERTS**, **ANALYTICS**, **USERS** categories for the widget to display. By default, the widget displays statistics information about all objects in the inventory. You can use the Select which tags to filter option to filter the information. The widget can interact with other widgets in the dashboard, taking data from them and displaying statistics . For example, you can have a Object List widget , which is the source of the data and an Environment Status widget, which is the destination. If you select objects and perform a multiselection interaction from the Object List widget, the Environment Status widget results are updated based on the selections you made in the Object List.

#### Where You Find the Environment Status Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Environment Status Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard.<br><br>The widget is also updated when it is in interaction mode. For example, when an item is selected in the provider widget, the content of the Environment Status widgets is refreshed. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Objects | The widget shows summarized information about the objects in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of resources. For example, if you select **Adapter Types › Container** from Select which tag to filter and click **Objects** and **Objects Collecting** , the widget displays the number of containers and collecting containers. |
| Metrics | The widget shows summarized information about available metrics. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of metrics. |
| Applications | The widget shows summarized information about available applications. You can filter the information that appears in self provider mode when you select an object from Select |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | which tag to filter. You can select what type of information to include in the summary of applications. |
| Alerts | The widget shows summarized information about alerts in your environment. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of alerts. |
| Analytics | The widget shows summarized information about the analytics plug-ins. You can filter the information that appears in self provider mode when you select an object from Select which tag to filter. You can select what type of information to include in the summary of analytics. |
| Users | The widget shows the number of users defined in VMware Aria OperationsVMware Cloud Foundation Operations. Select **Administration › Access Control › User Accounts** |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br><br>If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 5. To add another filter criteria set, click **Add another criteria set**. |

**Faults Widget**

The Faults widget displays detailed information about faults experienced by an object

The Faults widget configuration options are used to customize each instance of the widget that you add to your dashboards.

**Where You Find the Faults Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Faults Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

**Forensics Widget**

The Forensics widget shows how often a metric has a particular value as a percentage of all values, within a given time period. It can also compare percentages for two time periods.

**How the Forensics Widget and Configuration Options Work**

You can add the Forensics widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit the Forensics widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

**Where you Find the Forensics Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Forensics Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Percentile | Indicates how much data is above or below the specific value. For example, it indicates that 90% of the data is more than 4 when a vertical line occurs on the value 4. |
| **Input Data** | |
| | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** ico n.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list. |

## Geo Widget

If your configuration assigns values to the Geo Location object tag, the geo widget shows where your objects are located on a world map. The geo widget is similar to the **Geographical** tab on the Inventory page.

### How the Geo Widget and Configuration Options Work

You can move the map and zoom in or out by using the controls on the map. The icons at each location show the health of each object that has the Geo Location tag value. You can add the geo widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a Geo widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

### Where You Find the Geo Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Geo Widget Toolbar Options

| Option | Description |
|---|---|
| Zoom in | Zooms in on the map. |
| Zoom out | Zooms out on the map. |

### Geo Widget Configuration Options

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

### Heatmap Widget

The Heatmap widget contains graphical indicators that display the current value of two selected attributes of objects of tag values that you select. In most cases, you can select only from internally generated attributes that describe the general operation of the objects, such as health or the active anomaly count. When you select a single object, you can select any metric for that object.

## How the Heatmap Widget and Configuration Options Work

You can add the Heatmap widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The Heatmap widget has a General mode and an Instance mode. The General mode shows a colored rectangle for each selected resource. In the Instance mode, each rectangle represents a single instance of the selected metric for an object.

You can click a color or the size metric box in the bottom of the Heatmap widget to filter the display of cells in the widget. You can click and drag the color filter to select a range of colors. The Heatmap widget displays cells that match the range of colors.

When you point to a rectangle for an object, the widget shows the resource name, group-by values, the current values of the two tracked attributes, virtual machine details, the metric name, and the value of the color. Click **Show Sparkline** to view the value.

You edit a Heatmap widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

## Where You Find the Heatmap Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Heatmap Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Dashboard Navigation | Actions you can run on the selected alert.<br><br>For example, you use the option to open a vCenter, data center, virtual machine, or in the vSphere Web Client, allowing you to directly modify an object for which an alert was generated and fix any problems. |
| Group Zoom | You can roll-up non-significant resources with similar characteristics into groups to obtain only the relevant data among the thousands of resources in the system. The roll-up method improves performance and decreases the memory usage. The roll-up box encompasses the average color and the sum of the sizes of all the resources. You can view all the resources by zooming in the roll-up box. |
| Show/Hide Text | Show or hide the cell name on the heatmap rectangle. |
| Show Details | If you configure the Heatmap widget as a provider to another widget, such as the Metric Chart widget , you can double-click a rectangle to select that object for the widget. If the widget is in Metric mode, double-clicking a rectangle selects the resource associated with the metric and |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | provides that resource to the receiving widget. Optionally, you can select a cell from the heatmap and click the **Show Details** icon to see details about the cell. |
| Reset Interaction | Returns the widget to its initial configured state and undoes any interactions selected in a providing widget. |
| Reset Zoom | Resets the heatmap display to fit in the available space. |
| Heatmap Configuration Drop-down | Select from a list of predefined heatmaps. |

**Heatmap Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Configurations | List of saved heatmap configuration options. You can create a configuration and save it in the list. From the options on the right, you can also delete, clone, and reorder the configurations. |
| Name | Name of the widget. |
| Group by | First-level grouping of the objects in the heatmap. |
| Then by | Second-level grouping of the objects in the heatmap. |
| Relational Grouping | After you select the Group by and Then by objects, select the **Relational Grouping** check box to reorganize the grouping of the objects, and to relate the objects selected in the Group by text box with the objects selected in the Then by text box. |
| Mode | **General mode**<br>The widget shows a colored rectangle for each selected resource. The size of the rectangle indicates the value of one selected attribute. The color of the rectangle indicates the value of another selected attribute.<br>**Instance mode** |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Each rectangle represents a single instance of the selected metric for a resource. A resource can have multiple instances of the same metric. The rectangles are all the same size. The color of the rectangles varies based on the instance value. You can use instance mode only if you select a single resource kind. |
| Object Type | Object that is the basis for the widget data. |
| Size by | An attribute to set the size of the rectangle for each resource.<br><br>Resources that have higher values for the Size By attribute have larger areas of the widget display. You can also select fixed-size rectangles. In most cases, the attribute lists include only metrics that VMware Aria OperationsVMware Cloud Foundation Operations generates. If you select a resource kind, the list shows all the attributes that are defined for the resource kind. |
| Color by | An attribute to set the color of the rectangle for each resource. |
| Solid Coloring | Select this option to use solid colors instead of a color gradient. By default, the widget assigns red color for high value, brown color for intermediate value and green color for low value. Click the color box to set a different color for the values. You can add up to seven color thresholds by clicking color range. |
| Color | Shows the color range for high, intermediate and low values. You can set each color and type minimum and maximum color values in the **Min Value** and **Max Value** text boxes. By default, green indicates a low value and red indicates the high end of the value range. You can change the high and low values to any color and set the color to use for the midpoint of the range. You can also set the values to use for either end of the color range, or let VMware Aria OperationsVMware Cloud Foundation Operations define the colors based on the range of values for the attribute.<br><br>If you leave the text boxes blank, VMware Aria OperationsVMware Cloud Foundation Operations maps the highest and lowest values for the **Color By** metric to the end colors. If you set a minimum or maximum value, any metric at or beyond that value appears in the end color. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br><br>If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

### Health Widget

The Health widget is the status of the health-related alerts for the objects it is configured to monitor in VMware Aria OperationsVMware Cloud Foundation Operations. Health alerts usually require immediate attention. You can create one or more health widgets for different objects that you add to your custom dashboards.

#### How the Health Widget and Configuration Options Work

You can add the Health widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.

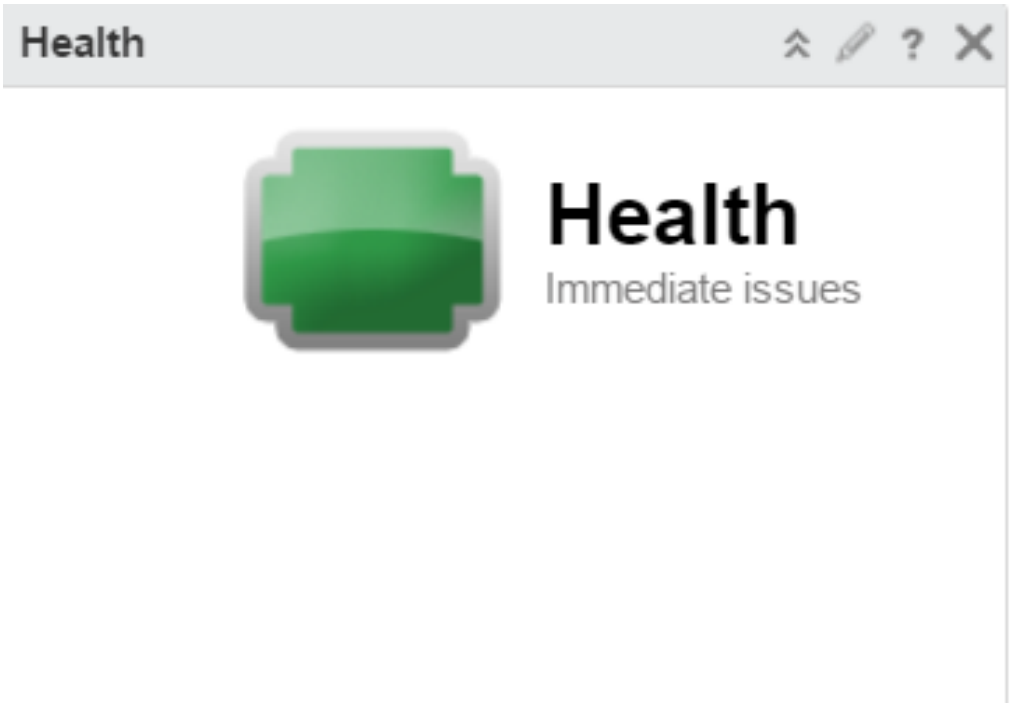


The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the **Badge Mode** configuration option is set to **Off**, the badge and a chart appears. The type of chart depends on the object that the widget is configured to monitor.

- A trend line displays the health status of the monitored object if the object does not provide its resources to any other object. For example, if the monitored object is a virtual machine or a distributed switch.
- A weather map displays the health of the ancestor and descendant objects of the monitored object for all other object types. For example, if the monitored object is a host that provides CPU and memory to a virtual machine.

If the **Badge Mode** is set to **On**, only the badge appears.

You edit a Health widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

**Where You Find the Health Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Heath Widget Display Options**

The Health widget displays a health badge. The widget also displays a health trend when not in badge mode.

| Option | Description |
|---|---|
| Health Badge | Status of the objects configured for this instance of the widget.<br><br>Click the badge to open the **Alerts** tab for the object that provides data to the widget.<br><br>If the **Badge Mode** option is off, a health weather map or trend chart appears for the object. Whether the map or chart appears depends on the object type. The health weather map displays tool tips for up to1000 objects. |
| Health Trend | Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** confi guration option is off. If the **Badge Mode** is on, only the badge appears. |

**Heath Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Badge Mode | Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.<br><br>Select one of the following options:<br>• On. Only the badge appears in the widget.<br>• Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

**Health Chart Widget**

The Health Chart widget displays Health, Risk, Efficiency, or custom metric charts for selected objects. You use the widget to compare the status of similar objects based on the same value or name.

**How the Health Chart Widget and Configuration Options Work**

You can add the Health Chart widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. The information that it displays depends on how the widget is configured.

If the widget is configured to display Health, Risk, or Efficiency, the chart values are based on the generated alerts for the selected alert type for the selected objects.

If the widget is configured to display custom metrics, chart values are based on the metric value for the configured time period.

You edit the Health Chart widget after you add it to the dashboard. The changes you make to the options create a custom widget with the selected charts.

The charts are based either on Health, Risk, or Efficiency alert status, or you can base them on a selected metric. You can include a single object, multiple objects, or all objects of a selected type.

To view the value of the object at a particular time, point your cursor over the chart. A date range and metric value tool tip appear.

A context drop-down menu for each chart can be accessed at the top-right corner after the last metric value.

For each chart, you can view the minimum, maximum, and last metric values. The values are displayed at the top-right corner of each chart. Each of the values is preceded by an appropriate icon of the same color as the state of the metric value.

If there is not enough space to view the metric values, a blue information icon is displayed. Point your cursor over the icon to view the metric value details.

### Where You Find the Health Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Health Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Date Controls | Use the date selector to limit the data that appears in each chart to the time period you are examining.<br><br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br><br>**Dashboard Time** is the default option. |

### Health Chart Widget Graph Selector Options

The graph selector options determine how individual data appears in the graph.

| Option | Description |
|---|---|
| Close | Deletes the chart. |
| Save a snapshot | Creates a PNG file of the current chart. The image is the size that appears on your screen.<br><br>You can retrieve the file in your browser's download folder. |
| Save a full screen snapshot | Downloads the current graph image as a full-page PNG file, which you can display or save.<br><br>You can retrieve the file in your browser's download folder. |
| Download comma-separated data | Creates a CSV file that includes the data in the current chart.<br><br>You can retrieve the file in your browser's download folder. |
| Units | Select the units in which the widget displays data. This option is visible when you select a custom source of data in the widget configuration. |

### Health Chart Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Order By | Determines how the object charts appear in the widget.<br><br>You can order them based on value or name, and in ascending or descending order. |
| Chart Height | Controls the height of all charts. Choose from three possible choices - Small, Medium, Large. Default is Medium. |
| Pagination number | Number of charts that appears on a page.<br><br>If you prefer scrolling through the charts, select a higher number. If you prefer to page through the results, select a lower number. |
| Auto Select First Row | Determines whether to start with the first row of data. |
| Metric | Determines the source of the data.<br><br>• Health, Risk, or Efficiency. The displayed charts are based on one of these alert badges.<br>• Custom. The displayed charts are based on the selected metric and use either alert symptom state colors or the selected custom color. You can select a unit for the custom metric from the drop-down menu or choose to allow the widget to automatically pick a unit. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | If you apply custom colors, enter the value in each box that is the highest or lowest value that should be that color. You can select a unit for the metric. |
| Metric Unit | Select a unit for the custom metric. |
| Show | Select one or more of the following items to display in the widget:<br>• Select **Object Name** to display the name of the object in the widget.<br>• Select **Metric Name** to display the name of the metric in the widget. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br>Click the **Select All** icon to select all the objects in the list.<br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br><br>If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

### Log Analysis Widget

The Log Analysis widget displays the logs for objects using VMware Aria Operations for Logs. You can configure the objects for which data appears in the widget, or provide objects using the dashboard widget interactions options.

The following screenshot shows the log analysis widget.

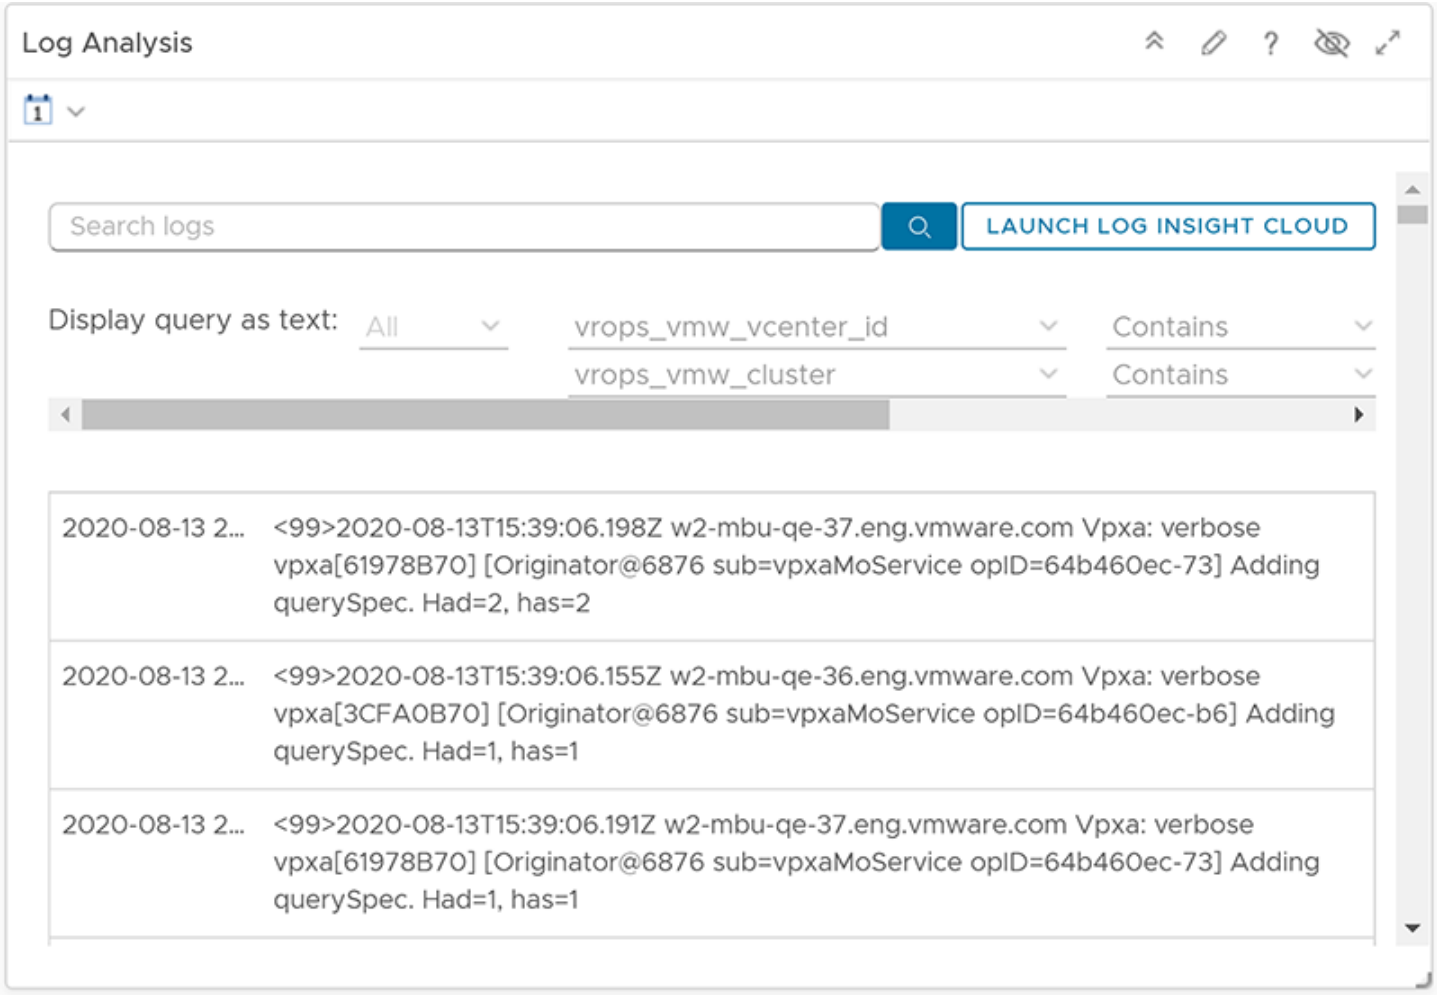


You can view, filter, and search the logs that are displayed. For in-depth analysis of the logs, you can run VMware Aria Operations for Logs.

## Where You Find the Log Analysis Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Log Analysis Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | • On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

**NOTE**
Do not use the log analysis form widget for reporting. The widget requires user inputs to show results. The widget is interactive and performs cross-product queries.

### Mashup Chart Widget

The Mashup Chart widget shows disparate pieces of information for a resource. It shows a health chart and metric graphs for key performance indicators (KPIs).

#### How the Mashup Chart Widget and Configuration Options Work

The Mashup Chart widget contains charts that show different aspects of the behavior of a selected resource. By default, the charts show data for the past six hours.

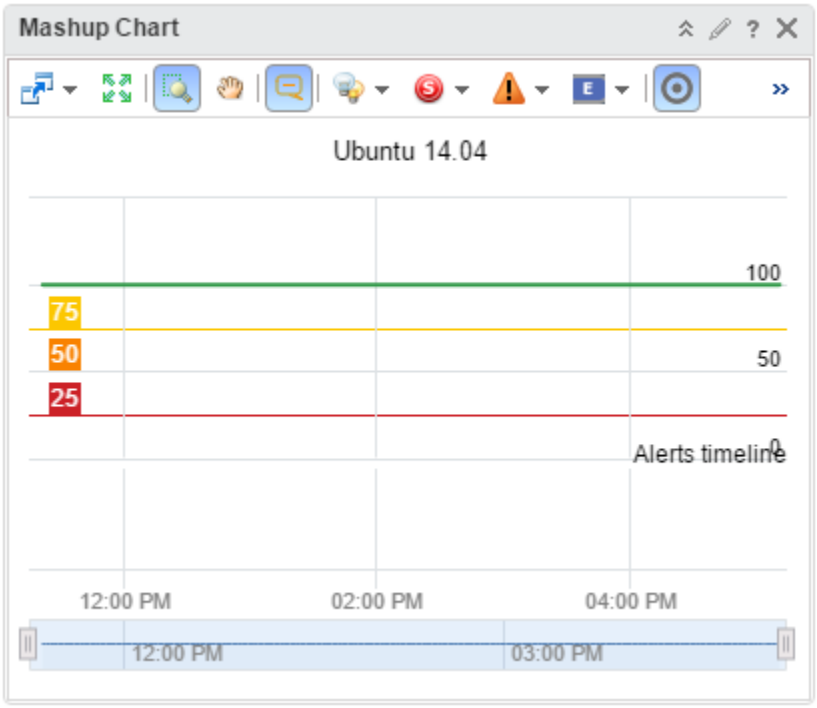


The Mashup Chart widget contains the following charts.

- A Health chart for the object, which can include each alert for the specified time period. Click an alert to see more information, or double-click an alert to open the Alert Summary page.
- Metric graphs for any or all the KPIs for any objects listed as a root cause object. For an application, this chart shows the application and any tiers that contain root causes. You can select the KPI to include by selecting **Chart Controls › KPIs** on the widget toolbar. Any shared area on a graph indicates that the KPI violated its threshold during that time period.

The metric graphs reflect up to five levels of resources, including the selected object and four child levels.

You edit a Mashup Chart widget after you add it to a dashboard. The changes you make to the options create a custom widget to meet the needs of the dashboard users.

## Where You Find the Mashup Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Mashup Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view.

| Option | Description |
|---|---|
| Filters | Filter data based on criticality, status, and alert type. |
| Event Filters | Filter based on the type of event such as, change, notification, and fault. |
| Date Controls | Use the date selector to limit the data that appears in each chart to the time period you are examining.<br><br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br><br>**Dashboard Time** is the default option. |
| Dashboard Navigation | You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. |

**Mashup Chart Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | object list and the **Tag Filter** pane to select an object based on tag values. |

**Metric Chart Widget**

You can use the Metric Chart widget to monitor the workload of your objects over time. The widget displays data based on the metrics that you select.

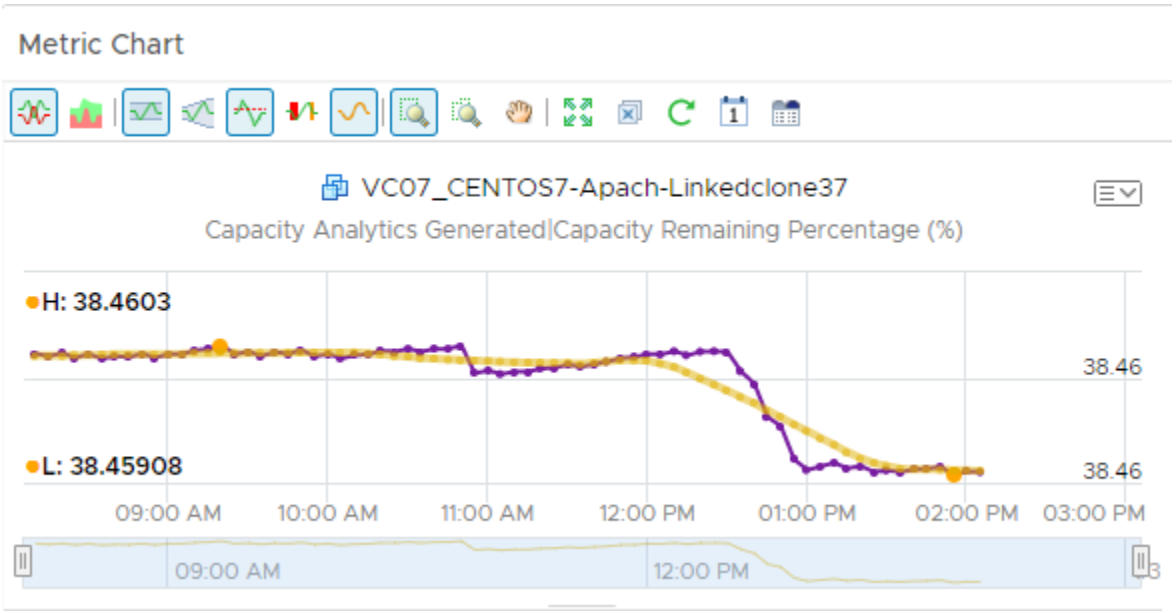


**How the Metric Chart Widget and Configuration Options Work**

You can add the Metric Chart widget to one or more custom dashboards and configure it to display the workload for your objects. The data that appears in the widget is based on the configured menu items for each widget instance.

You edit the Metric Chart widget after you add it to a dashboard. The changes you make to the menu items create a custom widget with the selected metrics that display the workload on your objects.

To select metrics, you can select an object from the object list, then select the metrics. Or, you can select a tag from the object tag list to limit the object list, then select an object. You can configure multiple charts for the same object or multiple charts for different objects.

To use the metric configuration, which displays a set of metrics that you defined in an XML file, the dashboard and widget configuration must meet the following criteria:

- The dashboard **Widget Interaction** menu items are configured so that another widget provides objects to the target widget. For example, an Object List widget provides the object interaction to a chart widget.
- The widget **Self Provider** options are set to **Off**.
- The custom XML file in the **Metric Configuration** drop-down menu is in the `/usr/lib/vmware-vcops/tools/opscli` directory and has been imported into the global storage using the import command.

### Where You Find the Metric Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Metric Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
|---|---|
| Split Charts | Displays each metric in a separate chart. |
| Stacked Chart | Consolidates all charts into one chart. This chart is useful for seeing how the total or sum of the metric values vary over time. To view the stacked chart, ensure that the split chart option is turned off. |
| Dynamic Thresholds | Shows or hides the calculated dynamic threshold values for a 24-hour period. |
| Show Entire Period Dynamic Thresholds | Shows or hides dynamic thresholds for the entire time period of the graph. |
| Static Thresholds | Shows or hides the threshold values that have been set for a single metric. |
| Anomalies | Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below. |
| Trend Line | Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points. |
| Show Data Values | Activates the data point tooltips if you switched to a zoom or pan option. **Show Data Point Tips** must be activated. |
| Zoom All Charts | Resizes all the charts that are open in the chart pane based on the area captured when you use the range selector.<br><br>You can switch between this option and **Zoom the View**. |
| Zoom the View | Resizes the current chart when you use the range selector. |
| Pan | When you are in zoom mode, allows you to drag the enlarged section of the chart so that you can view higher or lower, earlier or later values for the metric. |
| Zoom to Fit | Resets the chart to fit in the available space. |
| Remove All | Removes all the charts from the chart pane, allowing to you begin constructing a new set of charts. |
| Refresh Charts | Reloads the charts with current data. |
| Date Controls | Opens the date selector.<br><br>Use the date selector to limit the data that appears in each chart to the time period you are examining. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br><br>**Dashboard Time** is the default option. |
| Near Real-Time Monitoring | Displays near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days. |
| Generate Dashboard | Saves the current charts as a dashboard. |
| Save as PDF | Saves the current metric chart as a PDF file.<br>If you have added many metrics for different objects to a metric chart, you can quickly download and share the PDF file with another user or VM owner. You can create ad-hoc reports when analyzing metrics. You can also add notes to the metrics to provide context to the user. |

**Metric Chart Widget Graph Selector Options**

The graph selector options determine how individual data appears in the graph.

| Option | Description |
|---|---|
| Close | Deletes the chart. |
| Save a snapshot | Creates a PNG file of the current chart. The image is the size that appears on your screen.<br><br>You can retrieve the file in your browser's download folder. |
| Download comma-separated data | Creates a CSV file that includes the data in the current chart.<br><br>You can retrieve the file in your browser's download folder. |
| Save a full screen snapshot | Downloads the current graph image as a full-page PNG file, which you can display or save.<br><br>You can retrieve the file in your browser's download folder. |
| Units | You can display the data with dots or as a percentage. |
| Thresholds | You can choose to show/hide **Critical**, **Immediate**, and **Warning** thresholds in the current chart. |
| Scales | You can choose a scale for a stacked chart.<br>• Select **Linear** to view a chart in which the Y axis scale increases in a linear manner. For example, the Y axis can have ranges from 0 to 100, 100 to 200, 200 to 300, and so on.<br>• Select **Logarithmic** to view a chart in which the Y axis scale increases in a logarithmic manner. For example, the Y axis can have ranges from 10 to 20, 20 to 300, 300 to 4000, and so on. This scale gives a better visibility of minimum and maximum values in the chart when you have a large range of metric values.<br><br>**NOTE**<br>If you select a logarithmic scale, the chart does not display data points for metric values less than or equal to 0, which leads to gaps in the graph. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Select **Combined** to view overlapping graphs for the metrics. The chart uses individual scales for each graph instead of using a relative scale, and displays a combined view of the graphs.<br>• Select **Combined by Unit** to view a chart that groups the graphs for similar metric units together. The chart uses a common scale for the combined graphs. |
| Move Down | Moves the chart down one position. |
| Move Up | Moves the chart up one position. |

You can take the following actions on the Metric Chart graph.

| Option | Description |
|---|---|
| Y Axis | Shows or hides the Y-axis scale. |
| Chart | Shows or hides the line that connects the data points on the chart. |
| Data Point Tips | Shows or hides the data point tooltips when you hover the mouse over a data point in the chart. |
| Zoom by X | Enlarges the selected area on the X axis when you use the range selector in the chart to select a subset of the chart. You can use **Zoom by X** and **Zoom by Y** simult aneously. |
| Zoom by Y | Enlarges the selected area on the Y axis when you use the range selector in the chart to select a subset of the chart. You can use **Zoom by X** and **Zoom by Y** simult aneously. |
| Zoom by Dynamic Thresholds | Resizes the Y axis of the chart so that the highest and the lowest values on the axis are the highest and the lowest values of the dynamic threshold calculated for this metric. |
| Vertical resize | Resizes the height of a graph in the chart. |
| **Remove** icon next to each metric name in a stacked chart | Removes the graph for the metric from the chart. |

**Metric Chart Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br><br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br><br>1. Double-click a metric box in the list to customize the metric and click **Update**.<br><br>You can use the **Box Label** text box to customize the label of a metric box.<br><br>You can use the **Unit** text box to define a measurement unit of each metric.<br><br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br><br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br>2. Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Optionally, you can customize a metric and apply the customization to other metrics in the list.<br>1. Double-click a metric box in the list to customize the metric and click **Update**.<br>You can use the **Box Label** text box to customize the label of a metric box.<br>You can use the **Unit** text box to define a measurement unit of each metric.<br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br>2. Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

**Metric Picker Widget**

The Metric Picker widget displays a list of available metrics for a selected object.

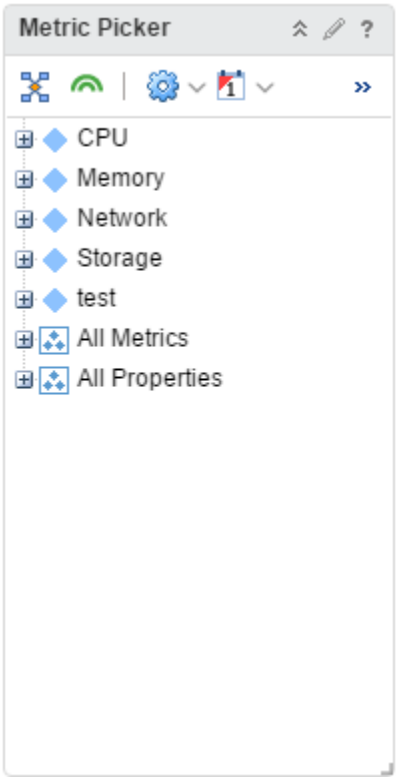


## How the Metric Picker Widget and Configuration Options Work

With the Metric Picker widget, you can check the list of the object's metrics. To select an object to pick its metrics, you use another widget as a source of data, for example, Topology Graph widget. To set a source widget that is on the same dashboard, you use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a dashboard that contains the source widget. You can also search for objects using tags.

You edit a Metric Picker widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

## Where You Find the Metric Picker Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Metric Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
|---|---|
| Show common metrics | Filter based on common metrics. |
| Show collecting metrics | Filter based on collecting metrics. |
| Metrics or Properties | Filter based on metrics or property metrics. |
| Time Range | Filter based on selected time range. |
| Search | Search for dashboards, views, and network IP addresses using tags. |

### Metric Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Action |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

## Object List Widget

The Object List widget displays a list of the objects available in the environment.

### How the Object List Widget and Configuration Options Work

The Object List widget displays a data grid with objects in the inventory. The default configuration of the data grid appears in Object List Widget Options section. You can customize it by adding or removing default columns. You can use the **Additional Column** option to add metrics when you configure the widget.

You edit an Object List widget after you add it to a dashboard. Configuration of the widget helps you observe parent and child objects. You can configure the widget to display the child objects of an object selected from another widget, for example, another Object List or Object Relationship widget, in the same dashboard.

Click the legend at the bottom of the widget to filter the objects based on threshold. Point your cursor over any of the boxes to view tooltips.

### Where You Find the Object List Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Object List Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Action | Selects from a set of actions specific for each object type. To see available actions, select an object from the list of objects and click the toolbar icon to select an action. For example, when you select a datastore object in the graph, you can select **Delete Unused Snapshots for Datastore**. |
| Dashboard Navigation | Navigates you to the object. For example, when you select a datastore from the list of objects and click **Dashboard Navigation**, you can open the datastore in vSphere Web Client. |
| Reset Grid Sort | Returns the list of resources to its original order. |
| Reset Interaction | Returns the widget to its initial configured state and undoes any interactions selected in a providing widget.<br><br>Interactions are usually between widgets in the same dashboard, or you can configure interactions between widgets on different dashboards. |
| Object Detail | Select an object and click this icon to show the Object Detail page for the object. |
| Perform Multi-Select Interaction | If the widget is a provider for another widget on the dashboard, you can select multiple rows and click this button. The receiving widget then displays only the data related to the selected interaction items.<br><br>Use Ctrl+click for Windows, or Cmd+click for Mac OS X, to select multiple individual objects or Shift+click to select a range of objects, and click the icon to activate the interaction. |
| Display Filtering Criteria | Displays the object information on which this widget is based. |
| Page Size | |
| Filter | Locate data in the widget.<br><br>You can search for objects or filter the list based on the values of the metrics or properties in the additional columns of the Configuration section. |

## Object List Widget Data Grid Options

The data grid provides a list of inventory objects on which you can sort and search.

| Option | Description |
|---|---|
| ID | Unique ID for each object in the inventory, randomly generated and produced by VMware Aria OperationsVMware Cloud Foundation Operations. |
| Name | Name of the object in the inventory. |
| Description | Displays the short description of the object given during creation of the object |
| Adapter Type | Shows the adapter type for each object. |
| Object Type | Displays the type of the object in the inventory. |
| Policy | Displays policies that are applied to the object. To see policy details and create policy configurations, in the menu click **Administration**, and then in the left pane click **Policies**. |
| Creation Time | Displays the date, time, and time zone of the creation of an object that was created in the inventory. |
| Identifier 1 | Can contain the custom name of the object in the inventory or default unique identifier, depending on the type of inventory object. For example, My_VM_1 for a VM in the inventory, or 64-bit hexadecimal value for VMware Aria OperationsVMware Cloud Foundation Operations Node. |
| Identifier 2 | Can contain the abbreviation of an object type and the unique decimal number or parent instance, depending on the type of the object. For example, vm-457 for a VM and an IP address for VMware Aria OperationsVMware Cloud Foundation Operations Node. |
| Identifier 3 | Can contain a unique number identifying an adapter type. For example, 64-bit hexadecimal value for vCenter Adapter |
| Identifier 4 | Additional unique identifiers for the object. This option varies and depends on the adapter type that the object uses. |
| Identifier 5 | Additional unique identifiers for the object. This option varies and depends on the adapter type that the object uses. |
| Object Flag | Displays a badge icon for each object. You can see the status when you point to the badge. |
| Collection State | Displays the collection state of an adapter instance of each object. You can see the name of the adapter instance and its state in a tool tip when you point to the state icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**. |
| Collection Status | Displays the collection status of the adapter instance of each object. You can see the name of the adapter instance and its status in a tool tip when you point to the status icon. To manage an adapter instance to start and stop collection of data, in the menu, click **Administration**, and then in the left pane click **Inventory**. |
| Relevance | Displays the user interest on objects based on the number of clicks. The relevance is determined using a system-wide |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | ranking algorithm that rates the object with most clicks as most relevant object. |
| Internal ID | Unique number that VMware Aria OperationsVMware Cloud Foundation Operations uses to identify the object internally. For example, the internal ID appears in log files used for troubleshooting. |

**Object List Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Auto Select First Row | Determines whether to start with the first row of data. |
| **Input Data** | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Filter**<br><br>**NOTE**<br>If there are more than ten objects under a section in the inventory tree, you can search for an object using the search option. If there are more than thousand objects in a section, use the **View More** button under the last object that is displayed to view the rest of the objects. | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied.<br><br>If the objects have an input transformation applied, you select tag values for the transformed objects. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |
| **Additional Columns** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | Add metrics based on object types. The selected metrics are displayed as additional columns in the widget.<br><br>1. Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br><br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br><br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2. Optionally, you can double-click a metric box in the list to customize the label of the metric and click **Update**. |

## Object Relationship Widget

The Object Relationship widget displays the hierarchy tree for the selected object. You can create one or more hierarchy trees in VMware Aria OperationsVMware Cloud Foundation Operations for the selected objects that you add to your custom dashboards.

### How the Object Relationship Widget and Configuration Options Work

You can add the Object Relationship widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

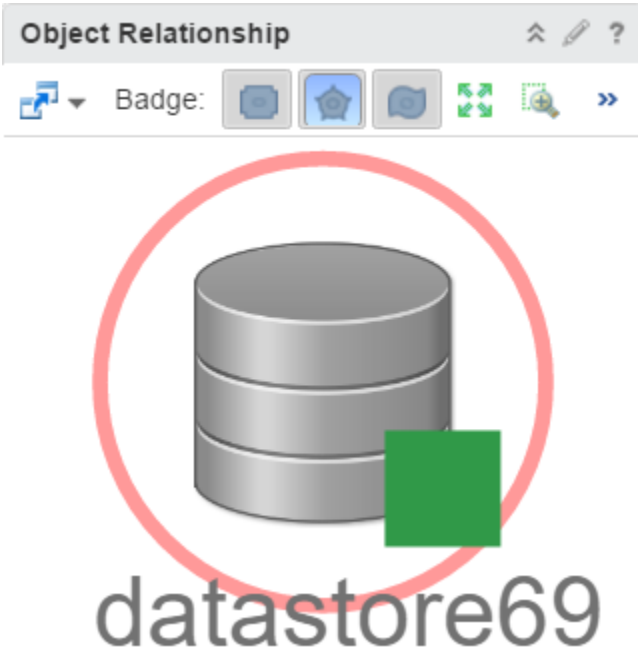


You edit an Object Relationship widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

### Where You Find the Object Relationship Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Object Relationship Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Dashboard Navigation | You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To be able to navigate to another dashboard, configure the relevant option when you create or edit the dashboard. |
| Badge | Displays the Health, Risk, or Efficiency alerts on the objects in the relationship map. You can select a badge for objects that appear in the widget. The tool tip of a badge shows the object name, object type, and the name of the selected badge with the value of the badge. You can only select one badge at a time. |
| Zoom to fit | Resets the chart to fit in the available space. |
| Pan | Click this icon and click and drag the hierarchy to show different parts of the hierarchy. |
| Show values on point | Shows or hides the data point tooltips when you hover the mouse over a data point in the chart. |
| Zoom the view | Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section. |
| Display Filtering Criteria | Shows the filtering settings for the widget in a pop-up window. |
| Zoom in | Zooms in on the hierarchy. |
| Zoom out | Zooms out on the hierarchy. |
| Reset to Initial Object | If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size. |
| Object Detail | Select an object and click this icon to show the Object Detail page for the object. |
| Show Alerts | Select the resource in the hierarchy and click this icon to show alerts for the resource. Alerts appear in a pop-up window. You can double-click an alert to view its Alert Summary page. |

**Object Relationship Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Auto Zoom to Fixed Node Size | You can configure a fixed zoom level for object icons in the widget display.<br><br>If your widget display contains many objects and you always need to use manual zooming, this feature is useful because you can use it to set the zoom level only once. |
| Node Size | You can set the fixed zoom level at which the object icons display. Enter the size of the icon in pixels.<br><br>The widget shows object icons at the pixel size that you configure. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

**Object Relationship (Advanced) Widget**

The Object Relationship (Advanced) widget displays a graph or tree view that depicts the parent-child relationship of the selected object. It provides advanced configuration options. You can create a graph or tree view in VMware Aria OperationsVMware Cloud Foundation Operations for the selected objects that you add to your custom dashboards.

**How the Object Relationship (Advanced) Widget and Configuration Options Work**

You can add the **Object Relationship (Advanced)** widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit an **Object Relationship (Advanced)** widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

You can double-click any object in the graph or tree view and see the specific parent-child objects for the focus object. When you double-click the object again, you see the original graph or tree view. If you point your cursor over an object icon, you see the health, risk, and efficiency details. You can also click the **Alerts** link for the number of generated alerts. Click the purple icon to view the child relationships of the object.

**Where You Find the Object Relationship (Advanced) Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Object Relationship (Advanced) Widget Toolbar Options**

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Options | Description |
|---|---|
| Dashboard Navigation | You can navigate to another dashboard when the object under consideration is also available in the dashboard to which you navigate. To navigate to another dashboard, configure the relevant option when you create or edit the dashboard. |
| Reset to Initial Object | If you change the hierarchy of the initial configuration or the widget interactions, click this icon to return to the initial resource. Clicking this icon also resets the initial display size. |
| Display Filtering Criteria | Shows the filtering settings for the widget in a pop-up window. |
| View Tree/View graph | Displays a tree or graph view of the relationships. |
| Vertical/Horizontal | Displays a vertical or horizontal view of the graph or tree view. |
| Hide Text/Show Text | Hides or displays the object names. |
| Standard View/Fit View | The **Standard View** option fixes the view to a specific zoom level<br><br>The **Fit View** option adjusts the graph or tree view to fit the screen. |
| Group Items/Ungroup Items | Groups by objects types. You can view further details by double-clicking on the object. You can also choose to display the graph or tree view without grouping the object types. |
| Path Exploration | Displays the relative relationship path between two selected objects on the graph or tree view. To highlight the path, click the **Path Exploration** icon and then select the two objects from the graph or tree view. |
| Layers | • **Parent/Child**: Displays a graph or tree view of the parent and child relationship for the specific object selected.<br>• **Custom**: Indicates the relationship between the objects that are part of the custom relationship. These objects have a connection via the selected custom relationship. |
| Quick Filter | Enter the name of an object that you want to see in the graph or tree view. |

**Object Relationship (Advanced) Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Name | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Parents Depth | Select the depth of parent objects to be displayed. |
| Children Depth | Select the depth of child objects to be displayed. |
| Inventory trees | Select an existing predefined traversal spec for the initial object relationship graph or tree view. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

**Property List Widget**

You can use the Property List widget to view the properties of objects and their values.

**How the Property List Widget and Configuration Options Work**

To observe the properties of objects in the Property List widget, you can select object property metrics when you configure the widget itself (Self Provider mode activated). Alternatively, you can select objects or object property metrics from another widget (Self Provider mode deactivated). You can also view a default or custom set of properties by selecting a preconfigured XML file in the Metric Configuration drop-down menu of the widget configuration window.

You edit a Property List widget after you add it to a dashboard. You can configure a widget to receive data from another widget by selecting **Off** for Self Provider mode. When the widget is not in Self Provider mode, it displays a set of predefined properties and their values of an object that you select on the source widget. For example, you can select a host on a Topology widget and observe its properties in the Property List widget. To configure the Property List as a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To configure a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

**Where You Find the Property List Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Property List Widget Data Grid Options**

The data grid provides information on which you can sort and search.

| Option | Description |
|---|---|
| Object Name | Name of the object, whose properties you observe. You can sort the properties by object name. To open the Object Details page, click an object name. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Property Name | Name of the property. You can sort the properties by property name. |
| Value | Value of the property. You can sort the properties by value. |

**Property List Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Visual Theme | Select a predefined visual style for each instance of the widget. The options are: Original and Compact. |
| Show Metric Full Name | You can choose to view the full name of the metrics. The options are: On and Off. |
| **Input Data** | |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br><br>You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**. |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | 1. Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br><br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br><br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2. Optionally, you can define measurement units for the metrics and properties in the list. Double-click a metric or properties box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**.<br>3. You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**. |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsection, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

**Recommended Actions Widget**

The Recommended Actions widget displays recommendations to solve problems in your vCenter instances. With recommendations, you can run actions on your data centers, clusters, hosts, and virtual machines.

**How the Recommended Actions Widget and Configuration Options Work**

The Recommended Actions widget appears on the Home dashboard, and displays the health status for the objects in your vCenter instance. At a glance, you can see how many objects are in a critical state, and how many objects need immediate attention.

From the Recommended Actions widget, you can focus in on problems further by, for example, clicking an object where the alerts triggered, and by clicking an individual alert.

You can edit the Recommended Actions widget on the Home dashboard, or on another dashboard where you add the widget. With the widget configuration options, you can assign a new name to the widget, set the refresh content, and set the refresh interval.

The Recommended Actions widget includes a selection bar, a summary pane, a toolbar for the data grid, and alert information for your objects in a data grid.

**Where You Find the Recommended Actions Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

**Recommended Actions Widget Selection Bar and Summary Pane**

| Option | Description |
|---|---|
| Scope | Allows you to select an instance of vCenter, and a data center in that instance. |
| Object tabs | Displays the object types with the number of objects affected in parentheses. You can display the actions for virtual machines, host systems, clusters, vCenter instances, and datastores. |
| Badge | Select the Health, Risk, or Efficiency badge to display alerts on your objects. Health alerts require immediate attention. Risk alerts require attention in the immediate future. Efficiency alerts require your input to reclaim wasted space or to improve the performance of your objects. For each badge, you can view critical, immediate, and warning alerts. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Health Status. With the Health badge selected, displays the number of affected objects and a summary of their health based on the alerts that triggered on the object. Lists the objects that have the worst health, and the number of alerts that triggered on each object.<br>• Risk Status. With the Risk badge selected, displays the number of affected objects and a summary of their risk based on the alerts that triggered on the object. Lists the objects that have the highest, and the number of alerts that triggered on each object.<br>• Efficiency Status. With the Efficiency badge selected, displays the number of affected objects. Lists the objects that have the lowest efficiency based on the alerts that triggered on the object, and the number of alerts that triggered on each object. |
| Search filter | Narrows the scope of the objects that appear. Enter a character or a number to search and display an object. When a filter is active, the name of the filter appears below the Search filter text box. |

**Recommended Actions Widget Toolbar Options**

The toolbar allows you to address an alert, and to filter the alert list.

| Option | Description |
|---|---|
| Cancel Alert | Cancels the selected alert.<br>You cancel alerts when you do not need to address them. Canceling the alert does not cancel the underlying condition that generated the alert. Canceling alerts is effective if the alert is generated by triggered fault and event symptoms because these symptoms are triggered again only when subsequent faults or events occur on the monitored objects. If the alert is generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again. |
| Suspend | Suspends an alert for a specified number of minutes.<br>You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object.<br><br>The user who suspends the alert becomes the assigned owner. |
| Quick Filter | Narrows the search to one of the available filter types. For example, you can display all alerts that are related to the Compliance Alert Subtype. |

**Recommended Actions Widget Data Grid Options**

The data grid displays the alerts that triggered on your objects. To resolve the problems indicated by the alerts, you can link to the alerts and the objects on which the alerts triggered.

For more information, see Accessing Alerts .

| Option | Description |
|---|---|
| Criticality | Criticality is the level of importance of the alert in your environment. The alert criticality appears in a tooltip when you hover the mouse over the criticality icon.<br><br>The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**. |
| Actionable | When an alert has an associated action, you can run the action on the object to resolve the alert. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Suggested Fix | Describes the recommendation to resolve the problem. For example, for Compliance alerts, the recommendation instructs you to use the *vSphere Hardening Guide* to resolve the problem. You can find the *vSphere Hardening Guides* at http://www.vmware.com/security/hardening-guides.html.<br>You can view other available recommendations and their associated actions, if any, to resolve the problem when you click the drop-down menu. |
| Name | Name of the object for which the alert was generated, and the object type, which appears in a tooltip when you hover the mouse over the object name.<br><br>Click the object name to view the object details tabs where you can begin to investigate any additional problems with the object. |
| Alert | Name of the alert definition that generated the alert.<br><br>Click the alert name to view the alert details tabs where you can begin troubleshooting the alert. |
| Alert Type | Describes the type of alert that triggered on the selected object, and helps you categorize the alerts so that you can assign certain types of alerts to specific system administrators. For example, Application, Virtualization/Hypervisor, Hardware, Storage, and Network. |
| Alert Subtype | Describes additional information about the type of alert that triggered on the selected object, and helps you categorize the alerts to a more detailed level than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration. |
| Time | Date and time that the alert triggered. |
| Alert ID | Unique identification for the alert. This column is hidden by default. |

## Recommended Actions Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

**Risk Widget**

The risk widget is the status of the risk-related alerts for the objects it is configured to monitor. Risk alerts in VMware Aria OperationsVMware Cloud Foundation Operations usually indicate that you should investigate problems in the near future. You can create one or more risk widgets for objects that you add to your custom dashboards.

**How the Risk Widget and Configuration Options Work**

You can add the risk widget to one or more custom dashboards and configure it to display data that is important to the dashboard users.

The state of the badge is based on your alert definitions. Click the badge to see the **Summary** tab for objects or groups configured in the widget. From the **Summary** tab, you can begin determining what caused the current state. If the widget is configured for an object that has descendants, you should also check the state of descendants. Child objects might have alerts that do not impact the parent.

If the Badge Mode configuration option is set to Off, the badge and a chart appear. The type of chart depends on the object type that the widget is configured to monitor.

- A population criticality chart displays the percentage of group members with critical, immediate, and warning risk alerts generated over time, if the monitored object is a group.
- A trend line displays the risk status of the monitored object for all other object types.

If the Badge Mode is set to On, only the badge appears.

You edit a risk widget after you add it to a dashboard. The changes you make to the options create a custom widget that provides information about an individual object, a custom group of objects, or all the objects in your environment.

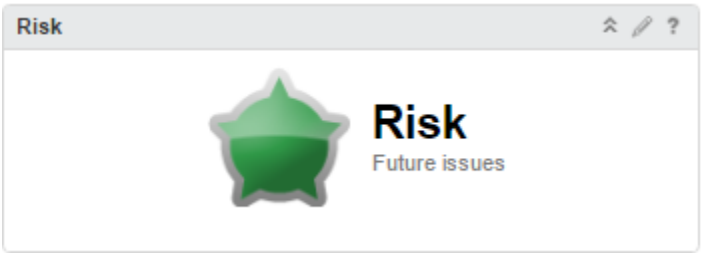


**Where You Find the Risk Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Risk Widget Display Options

The Risk Widget displays a risk badge. The widget also displays a risk trend chart when not in badge mode.

| Option | Description |
|---|---|
| Risk Badge | Status of the objects configured for this instance of the widget.<br><br>Click the badge to open the **Alerts** tab for the object that provides data to the widget. |
| Risk Trend | Displays a chart, depending on the selected or configured object. The charts vary, depending on whether the monitored object is a group, a descendent object, or an object that provides resources to other objects. The chart appears only if the **Badge Mode** configuration option is off. If the **Badge Mode** is on, only the badge appears. |

### Risk Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Badge Mode | Determines whether the widget displays only the badge, or the badge and a weather map or trend chart.<br><br>Select one of the following options:<br>• On. Only the badge appears in the widget.<br>• Off. The badge and a chart appear in the widget. The chart provides additional information about the state of the object. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

### Rolling View Chart Widget

The Rolling View Chart widget cycles through selected metrics at an interval that you define and shows one metric graph at a time. Miniature graphs, which you can expand, appear for all selected metrics at the bottom of the widget.

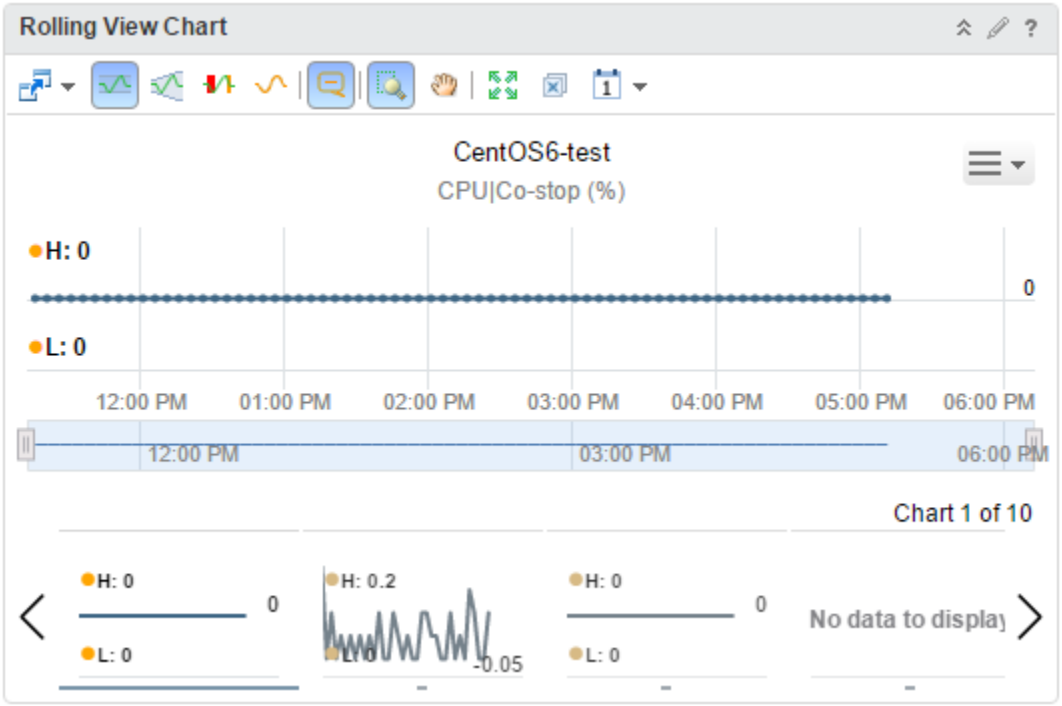


### How the Rolling View Chart Widget and Configuration Options Work

The Rolling View Chart widget shows a full chart for one selected metric at a time. Miniature graphs for the other selected metrics appear at the bottom of the widget. You can click a miniature graph to see the full graph for that metric, or set the

widget to rotate through all selected metrics at an interval that you define. The key in the graph indicates the maximum and minimum points on the line chart.

You edit a Rolling View Chart widget after you add it to a dashboard. The changes you make to the options create a custom chart to meet the needs of the dashboard users.

## Where You Find the Rolling View Chart Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Rolling View Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
|---|---|
| **Trend Line** | Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points. |
| **Dynamic Thresholds** | Shows or hides the calculated dynamic threshold values for a 24-hour period. |
| **Show Entire Period Dynamic Thresholds** | Shows or hides dynamic thresholds for the entire time period of the graph. |
| **Anomalies** | Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below. |
| **Zoom to Fit** | Changes all graphs to show the entire time period and value range. |
| **Zoom the view** | Click this icon and drag to outline a part of the hierarchy. The display zooms to show only the outlined section. |
| **Pan** | Click this icon and click and drag the hierarchy to show different parts of the hierarchy. |
| **Show Data Values** | After you click the **Show data point tips** icon to retrieve the data, click this icon and point to a graphed data point to show its time and exact value. In non-split mode, you can hover over a metric in the legend to show the full metric name, the names of the adapter instances (if any) that provide data for the resource to which the metric belongs, the current value, and the normal range. If the metric is currently alarming, the text color in the legend changes to yellow or red, depending on your color scheme. Click a metric in the legend to highlight the metric in the display. Clicking the metric again toggles its highlighted state. |
| **Date Controls** | Use the date selector to limit the data that appears in each chart to the time period you are examining. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br><br>**Dashboard Time** is the default option. |

**Rolling View Chart Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Auto Transition Interval | Time interval for a switch between charts in the widget. |
| **Input Data** | |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** ico n.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br><br>You can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**. |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.<br><br>1. Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br><br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br><br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2. Optionally, you can define measurement units for the metrics in the list. Double-click a metric box in the list, select a measurement unit in the **Unit** drop-down menu, and click **Update**. |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

**Scoreboard Widget**

The Scoreboard widget shows the current value for each metric of objects that you select.

**How the Scoreboard Widget and Configuration Options Work**

Each metric appears in a separate box. The value of the metric determines the color of the box. You define the ranges for each color when you edit the widget. You can customize the widget to use a sparkline chart to show the trend of changes of each metric. If you point to a box, the widget shows the source object and metric data. Icons in the box indicate the level of criticality.

You edit a Scoreboard widget after you add it to a dashboard. The widget can display metrics of the objects selected during editing of the widget or selected on another widget. When the Scoreboard widget is not in Self Provider mode, it shows metrics defined in a configuration XML file that you select in the Metric Configuration. It shows 10 predefined metrics if you do not select an XML file or if the type of the selected object is not defined in the XML file.

For example, you can configure the Scoreboard widget to use the sample Scoreboard metric configuration and to receive objects from the Topology Graph widget. When you select a host on a Topology Graph widget, the Scoreboard widget shows the workload, memory, and CPU usage of the host.

To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

**Where You Find the Scoreboard Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the

dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Scoreboard Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options.<br><br>When the Scoreboard widget is not in self-provider mode, it shows metrics defined in a configuration XML file that you select in the Metric Configuration. |
| Round Decimals | Select the number of decimal places to round the scores that the widget displays. |
| Box Columns | Select the number of columns that appear in the widget. |
| Layout Mode | Select a Fixed Size or Fixed View layout. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Fixed Size<br><br>Fixed View | Use these options to customize the size of the box for each object. |
| Old metric values | Select **Show** if you want the widget to show the previous value of the metric, if the current value is not available. Select **Hide** to hide the previous value of the metric, if the current value is not available. |
| Visual Theme | Select a predefined visual style for each instance of the widget. |
| Max Scores Count | Use these menus to customize the format of the scores that the widget displays. |
| Show | Select one or more of the following items to display in the widget:<br>• Select **Object Name** to display the name of the object in the widget.<br>• Select **Metric Name** to display the name of the metric in the widget.<br>• Select **Metric Unit** to display the metric unit in the widget.<br>• Select **Sparkline** to display the Sparkline chart for each metric. |
| Period Length | Select a length of time for the statistic information that the sparkline chart displays. |
| Show DT | Select an option to show or hide the dynamic threshold for the sparkline chart. |
| **Input Data** | |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** ico n.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br><br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br><br>1. Double-click a metric box in the list to customize the metric and click **Update**.<br><br>You can use the **Box Label** text box to customize the label of a metric box.<br><br>You can use the **Unit** text box to define a measurement unit of each metric.<br><br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **O range**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br><br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br><br>You can use the **Link to** option to add links to external and internal pages. Internal links open in the same tab. External links open in a new tab. Examples of external links are URLs whose hostname does not match with the current VMware Aria OperationsVMware Cloud Foundation Operations instance hostname. Internal links are URLs whose hostname matches the current VMware Aria OperationsVMware Cloud Foundation Operations instance hostname or starts with `index.action`.<br>2. Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.<br><br>1. Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br><br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
| --- | --- |
| | For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br><br>1. Double-click a metric box in the list to customize the metric and click **Update**.<br><br>You can use the **Box Label** text box to customize the label of a metric box.<br><br>You can use the **Unit** text box to define a measurement unit of each metric.<br><br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **O range**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br><br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **Color Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **Orange**, and 10 for **Red**.<br><br>You can use the **Link to** option to add links to external and internal pages. Internal links open in the same tab. External links will open in a new tab. Examples of external links are URLs whose hostname does not match with the current VMware Aria OperationsVMware Cloud Foundation Operations instance hostname. Internal links are URLs whose hostname matches the current VMware Aria OperationsVMware Cloud |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Foundation Operations instance hostname or starts with *index.action*.<br>2. Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

Click the **Near Real-Time Monitoring** icon to display near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days.

### Scoreboard Health Widget

The Scoreboard Health widget displays color-coded health, risk, efficiency, and custom metrics scores for objects that you select.

#### How the Scoreboard Health Widget and Configuration Options Work

The icons for each object are color coded to give a quick indication of the state of the object. You can configure the widget to display the scores of common or specific metrics of the object. You can use the symptom state color code or you can define your criteria to color the images. If you configure the widget to show the metric for objects that do not have this metric, those objects have blue icons.

You can double-click an object icon to show the Object Detail page for the object. When you point to the icon, a tool tip shows the name of the object and the name of the metric.

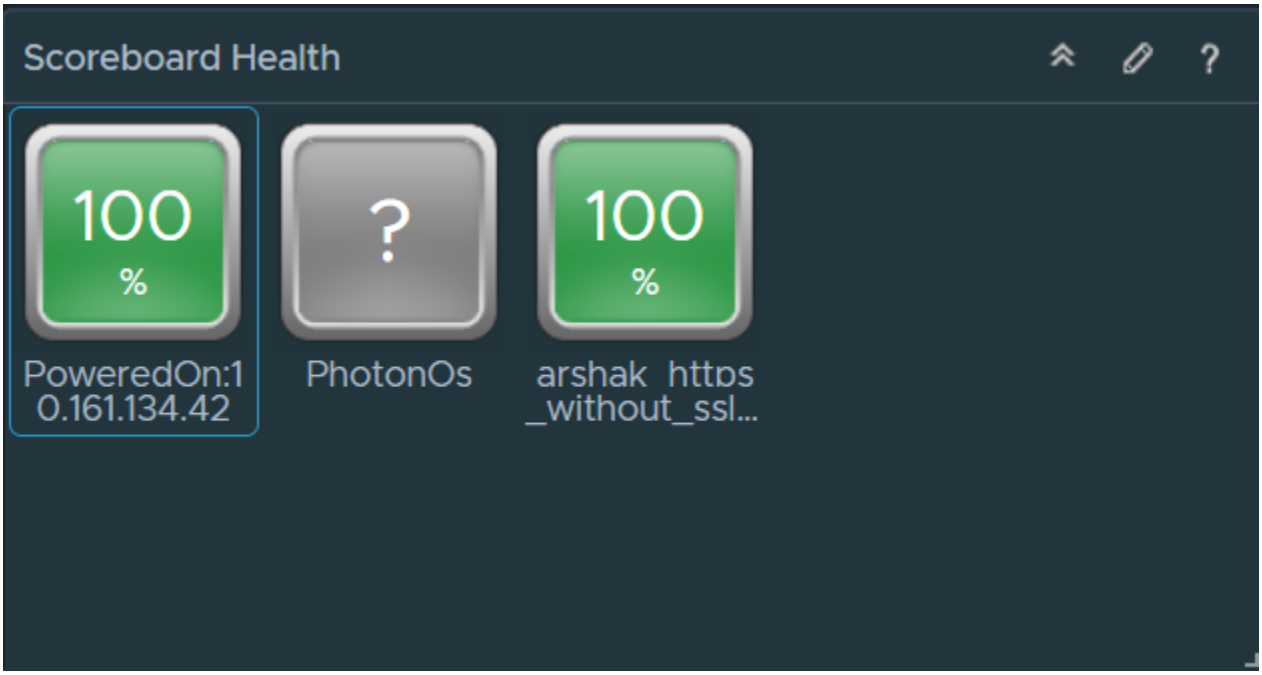


You edit a Scoreboard Health widget after you add it to a dashboard. To configure the widget, click the pencil at the upper-right corner of the widget window. The widget can display metrics of the objects that you select when you edit the widget, or that you select on another widget. For example, you can configure the widget to show the CPU workload of an object that you select on the Topology Graph widget. To set a source widget that is on the same dashboard, you must use the Widget Interactions menu when you edit a dashboard. To set a source widget that is on another dashboard, you must use the Dashboard Navigation menu when you edit the source dashboard.

## Where You Find the Scoreboard Health Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Scoreboard Health Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The **Configuration** section provides general configuration options for the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Image Type | Select an image type for the metrics. |
| Metric | Select the default or custom metric. |
| Pick Metric | Active only when you select **Custom** from the **Metric** menu.<br>Use to select a custom metric for the objects that the widget displays. Click **Pick Metric** and select an object type from the Object Type pane.<br><br>Use the Metric Picker pane to select a metric from the metric tree and click **Select Object** to check the objects from the type that you select on the Object Types pane. |
| Use Symptom state to color chart | Select to use the default criteria to color the image. |
| Custom ranges | Use to define custom criteria to color the image. You can define a range for each color. |
| **Input Data** | |
| | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |

### Sparkline Chart Widget

The Sparkline Chart widget displays graphs that contain metrics for an object in VMware Aria OperationsVMware Cloud Foundation Operations. You can use VMware Aria OperationsVMware Cloud Foundation Operations to create one or more graphs that contain metrics for objects that you add to your custom dashboards.

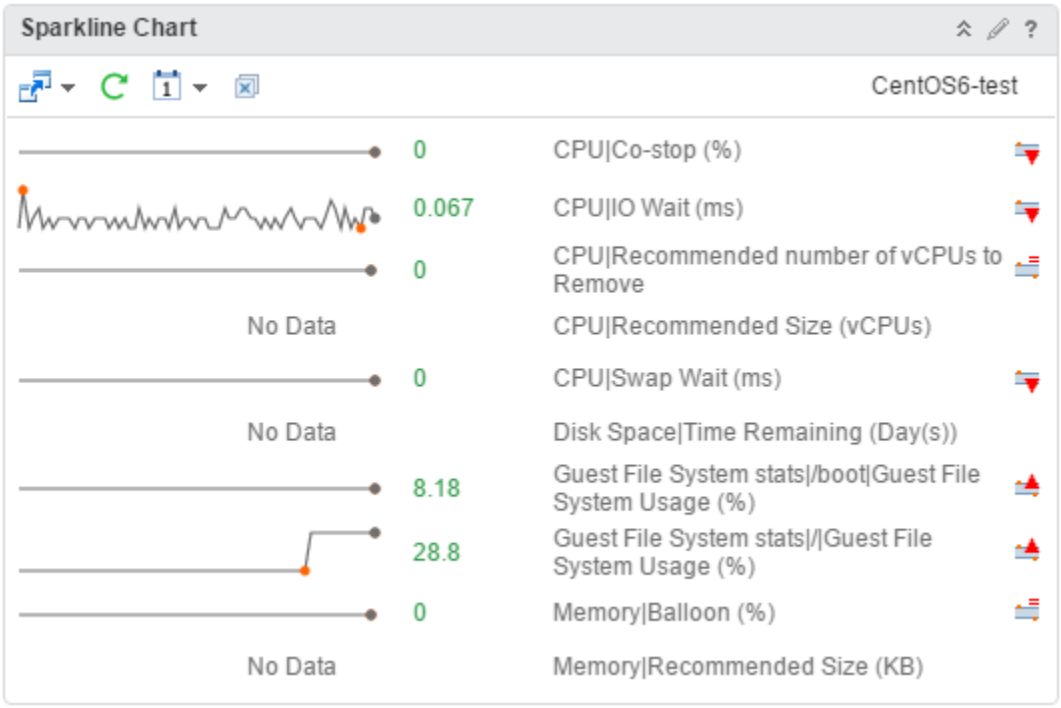


#### How the Sparkline Chart Widget and Configurations Options Work

If the metrics in the Sparkline Chart are for an object that another widget provides, the object name appears at the top right of the widget. If you select a metric when you edit the widget configuration, the widget uses the metric and its corresponding object as the source for dashboard interactions. The line in the graphs represents the average value of the selected metric for the specified time period. The boxed area in the graph represents the dynamic threshold of the metric.

Point to a graph in the Sparkline Chart widget to view the value of a metric in the form of a tool tip. You can also view the maximum and minimum values on a graph. The values are displayed as orange dots.

You can add the Sparkline Chart widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

The metrics shown in sparkline widget is the current value, to view the average value you can use transformations in list views or distribution charts to calculate an average. Another way to get to an average value is to double click on the sparkline to open the metric chart, click and drag to select a range, keep the mouse button depressed and hover for a few seconds, you should see a popup that has average value.

## Where You Find the Sparkline Chart Widget

The widget might be included on any of your custom dashboards. On the menu, click **Dashboards** to display a list of dashboards in the left pane.

## Sparkline Chart Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Option | Description |
|---|---|
| Dashboard Navigation | You can navigate to another dashboard when the object you select is also available in the dashboard to which you want to navigate. |
| Refresh | Refreshes the widget data. |
| Time Range | Select the range for the time period to show on the graphs. You can select a period from the default time range list or select start and end dates and times.<br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours.<br><br>**Dashboard Time** is the default option. |
| Near Real-Time Monitoring | Displays near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days. |
| Remove All | Removes all graphs. |

## Sparkline Chart Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Show Object Name | You can view the name of the object before the metric name in the Sparkline Chart widget.<br>• **On**. Displays the name of the object before the metric name in the widget.<br>• **Off**. Does not display the name of the object in the widget. |
| Column Sequence | Select the order in which to display the information.<br>• **Graph First**. The metric graph appears in the first column in the widget display.<br>• **Label First**. The metric label appears in the first column in the widget display. |
| Show DT | Select an option to show or hide the dynamic threshold for the sparkline chart. |
| **Input Data** | |
| Metrics | Select metrics on which you want to base the widget data. You can select an object and pick its metrics.<br>1. Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** ico n.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select metrics from the list and click the **Remove Selected Metrics** icon to remove the selected metrics.<br><br>Click the **Select All** icon to select all the metrics in the list.<br><br>Click the **Clear Selection** icon to clear your selection of metrics in the list.<br><br>Optionally, you can customize a metric and apply the customization to other metrics in the list.<br><br>1. Double-click a metric box in the list to customize the metric and click **Update**.<br><br>You can use the **Box Label** text box to customize the label of a metric box.<br><br>You can use the **Unit** text box to define a measurement unit of each metric.<br><br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **O range**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br><br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **C olor Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **O range**, and 10 for **Red**.<br>2. Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** chec k box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| Empty drop-down menu | Specifies a list with attributes to display. |
| | Add metrics based on object types. The objects corresponding to the selected metrics are the basis for the widget data.<br><br>Click the **Add New Metrics** icon to add metrics for the widget data. Select an object to view its metric tree and pick metrics for the object. The picked metrics appear in a list in this section.<br><br>The metric tree shows common metrics for several objects when you click the **Show common metrics** icon.<br><br>While selecting objects for which you want to pick metrics, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Optionally, you can customize a metric and apply the customization to other metrics in the list.<br><br>1. Double-click a metric box in the list to customize the metric and click **Update**.<br><br>You can use the **Box Label** text box to customize the label of a metric box.<br><br>You can use the **Unit** text box to define a measurement unit of each metric.<br><br>You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **O range**, and **Red** text boxes. You can also set coloring by symptom definition. If you do not want to use color, select **None**.<br><br>For example, to view the remaining memory capacity of a VM, select **Virtual Machine** as an object type, expand the **Memory** from the metric tree and double-click **Capacity Remaining(%)**. Define a meaningful label name and measurement unit to help you when you observe the metrics. You can select **Custom** from the **C olor Method** drop-down menu and specify different values for each color, for example 50 for **Yellow**, 20 for **O range**, and 10 for **Red**.<br>2. Select a metric and click the **Apply to All** icon to apply the customization for the selected metric to all the metrics in the list. |
| **Output Filter** | |
| | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 5. To add another filter criteria set, click **Add another criteria set**. |

### Tag Picker Widget

The Tag Picker widget lists all available object tags.

#### How the Tag Picker Widget and Configuration Options Work

With the Tag Picker widget, you can check the list of the object tags. You can use the widget to filter the information that another widget shows. You can select one or more tags from the object tree or search for tags, and the destination widget displays information about the objects with this tag. For example, you can select **Object Types › Virtual Machine** on the Tag Picker widget to observe statistic information about the VMs on the Environment Status widget.

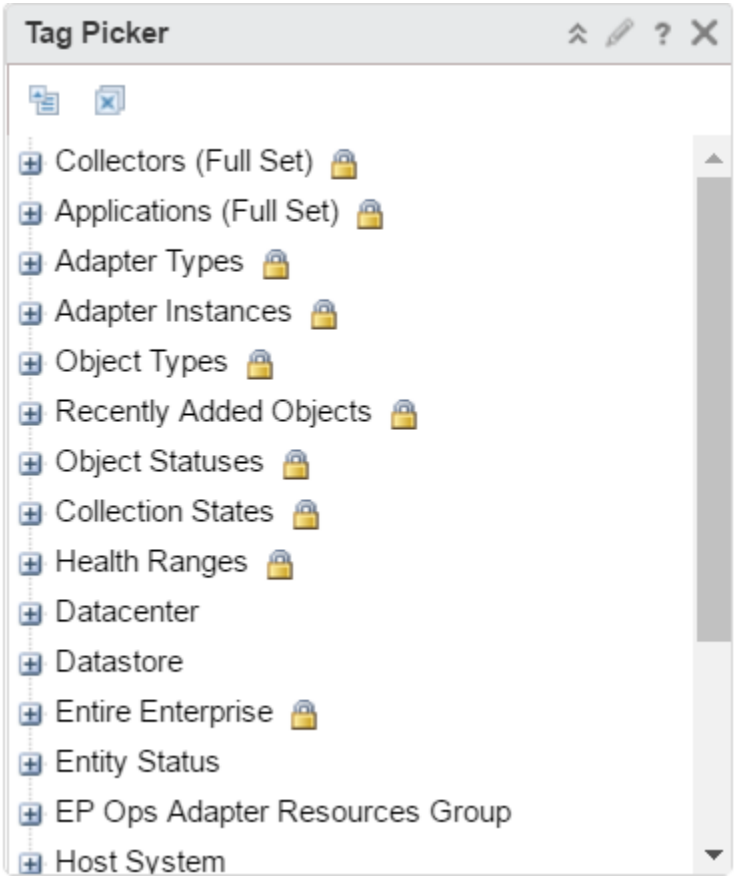


You edit a Tag Picker widget after you add it to a dashboard. To configure the widget, click the pencil in the upper right of the widget window. You can configure the Tag Picker widget to send information to another widget on the same dashboard or on another dashboard. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard. You can configure two Tag Picker widgets to interact when they are on different dashboards.

### Where You Find the Tag Picker Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Tag Picker Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Collapse All | Close all expanded tags and tag values. |
| Deselect All | Remove all filtering and view all objects in the widget. |
| Tag Picker | Select an object from your environment. |
| Dashboard Navigation | **NOTE**<br>Appears on the source widget and when the destination widget is on another dashboard.<br><br>Use to explore the information on another dashboard. |

### Tag Picker Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |

### Text Display Widget

You can use the Text Display widget to show text in the user interface. The text appears in the Text Display widget on the dashboard.

The Text Display widget can read text from a Web page or text file. You specify the URL of the Web page or the name of the text file when you configure the Text widget. To use the Text Display widget to read text files you must set a property in the *web.properties* file to specify the root folder that contains the file.

You can enter content in the Text Display widget in plain text or rich text format based on the view mode that you configure. Configure the Text Display widget in HTML view mode to display content in rich text format. Configure the Text Display widget in Text mode to display content in plain text format.

The Text Display widget can display websites that use the HTTPS protocol. The behavior of the Text Display widget with websites that use HTTP, depends on the individual settings of the websites.

**NOTE**
If the webpage that you are linking to has `X-Frame-Options` set to `sameorigin`, which denies rendering a page in an iframe, the Text Display widget cannot display the contents of the webpage.

#### How the Text Display Widget Configuration Options Work

You can configure the widget in the Text view mode or HTML view mode. In the HTML view mode, you can click **Edit** in the widget and use the rich text editor to add content.

If you configure the widget to use Text view mode, you can specify the path to the directory that contains the files to read or you can provide a URL. The content in the URL will be shown as text. If you do not specify a URL or text file, you can add content in the widget. Double-click the widget and enter content in plain text.

You can also use command-line interface (CLI) commands to add file content to the Text Display widget.

- To view a list of parameters, run the **`file -h|import|export|delete|list txtwidget`** command.
- To import text or HTML content, run the **`import txtwidget input-file [--title title] [--force]`** command.
- To export the content to the file, run the **`export txtwidget all|title[{,title}] [output-dir]`** command.
- To delete imported content, run the **`delete txtwidget all|title[{,title}]`** command.
- To view the titles of the content, run the **`list txtwidget`** command.

#### Where You Find the Text Display Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the

dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Text Display Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| View mode | Display text in text or rich text format. You can configure the widget in HTML view mode only when the **URL** and **File** fields are blank. |
| URL | Enter the URL. |
| File | Navigate to the file that contains the source text file by clicking the **Select** button.<br><br>To add, edit, and remove source text files, go to the **Text Widget Content** tile in the Configuration page. From the left menu, click **Operations › Configurations**, and then click the **Text Widget Content**tile from the VMware Aria OperationsVMware Cloud Foundation Operations user interface. |
| Test | Validates the correctness of the text file or URL that you enter. |

### Time Remaining Widget

The Time Remaining widget displays how much time remains before the resources of the object are exhausted.

VMware Aria OperationsVMware Cloud Foundation Operations calculates the percentage by object type based on historical data for the pattern of use for the object type. You can use the time remaining percentage to plan provisioning of physical or virtual resources for the object or rebalance the workload in your virtual infrastructure.

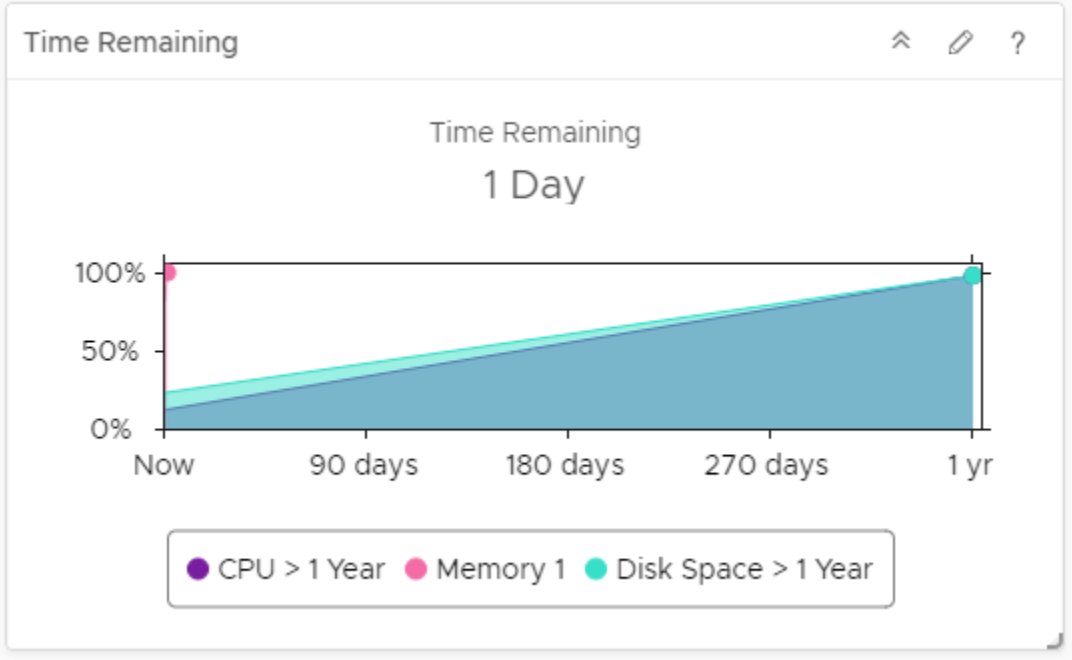


#### Where You Find the Time Remaining Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

#### Time Remaining Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

### Top Alerts Widget

Top alerts are the alerts with the greatest significance on the objects it is configured to monitor in VMware Aria OperationsVMware Cloud Foundation Operations. These are the alerts most likely to negatively affect your environment and you should evaluate and address them.

#### How the Top Alerts Widget and Configuration Options Work

You can add the top alerts widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

You edit a top alerts widget after you add it to a dashboard. The changes you make to the options help create a custom widget to meet the needs of the dashboard users.

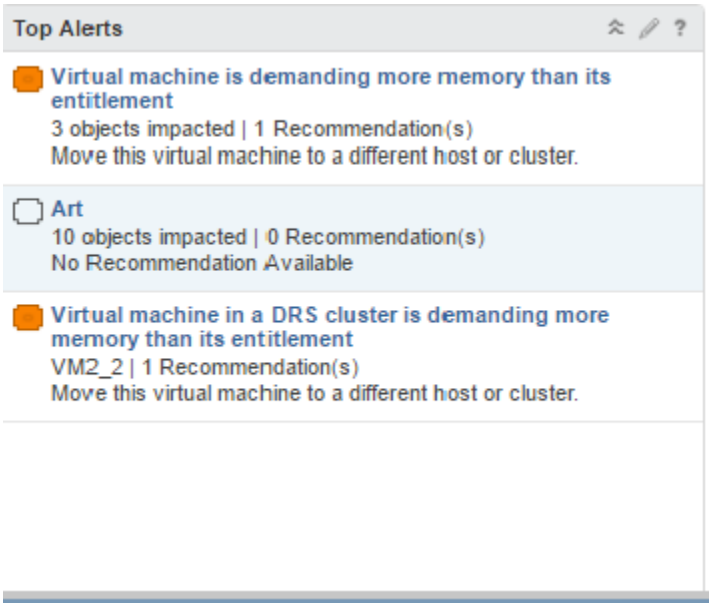


## Where You Find the Top Alerts Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

## Top Alerts Widget Display Options

The Top Alerts widget includes the short description of alerts configured for the widget. The alert name opens a secondary window from which you can link to the alert details. In the alert details, you can begin resolving the alerts.

| Option | Description |
|---|---|
| Alert name | Name of the generated alert. Click the name to open the alert details. |
| Alert description | Number of affected objects, and the number of recommendations and the best recommendation to resolve the alert. |

## Top Alerts Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Impact Badge | Select the badge for which you want alerts to appear.<br><br>The affected badge is configured when you configure the alert definition. |
| Number of Alerts | Select the maximum number of alerts to display in the widget. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |

### Top-N Widget

The Top-N widget displays the top n results from analysis of an object or objects that you select.

#### How the Top-N Widget and Configuration Options Work

You can select an object when you configure the Top-N widget or you can select an object on another widget. The widget shows an analysis of the applications, alerts, and metrics of an object and its child objects depending on how you configure the widget. The widget can show an analysis of the current values or values over a period of time. You can

receive detailed information about each object on the widget. When you double-click an object, the Object Detail page appears.

You can configure a widget to receive data from another widget by selecting **Off** for Self Provider. You can configure a widget to display results from analysis of an object that you select on the source widget.

For example, you can select a host on a Topology widget and observe the metric analysis of the virtual machines on the host. To set a receiver widget that is on the same dashboard, use the **Widget Interactions** menu when you edit a dashboard. To set a receiver widget that is on another dashboard, use the **Dashboard Navigation** menu when you edit a source dashboard.

### Where You Find the Top-N Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Top-N Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains icons that you can use to change the view of the graphs.

| Icon | Description |
|---|---|
| Dashboard Navigation | Takes you to a predefined object. For example, when you select a datastore from the data grid and click **Dashboard Navigation**, you can open the datastore in the vSphere Web Client. |
| Select Date Range | Limits the alerts that appear in the list to the selected date range.<br><br>Select **Dashboard Time** to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. |
| Object details | Select an object and click this icon to show the Object Detail page for the object. |
| Display Filtering Criteria | Shows the filtering settings for the widget in a pop-up window. |
| Percentile | Filters and displays objects based on the percentile entered while configuring the widget. You can change the percentile using the drop-down option.<br><br>**NOTE**<br>The Percentile option is activated in the toolbar of the widget after you configure the widget with the Percentile option. To configure the widget with the Percentile option, navigate to **Configurations › Top-N Options › Metric Analysis › Percentile**. |

### Top-N Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

The **Input Transformation** section provides options to transform the input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

The **Additional Columns** section provides options to select metrics that are displayed as additional columns in the widget.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Redraw Rate | Set the redraw rate. |
| Bars Count | Select the number of top results. |
| Round Decimals | Select the number of decimals to round the scores displayed in the widget. |
| Filter old metrics | Select or deselect whether the analysis includes old metric values. |
| Application Health and Performance | • Top Least Healthy. The top n results from an analysis of the object or objects that are the least healthy.<br>• Top Most Healthy. The top n results from an analysis of the object or objects that are the most healthy.<br>• Top Most Volatile. The sorted list of values based on the standard deviation of values for several alerts over time.<br>Select the criteria for analysis of the objects. |
| Alert Analysis | Select the criteria for analysis of the alerts. |
| Metric Analysis | If you select this option, you must select a metric in the **Output Data** section. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Top Highest Utilization. A list of objects with similar object types that have the highest utilization on configuring usage metrics like CPU usage and memory usage.<br>• Top Lowest Utilization. A list of objects with similar object types that have the lowest utilization on configuring usage metrics like CPU usage and memory usage.<br>• Top Abnormal States. The objects are ordered by the duration of all alarms that are triggered on the selected metric for a selected interval.<br>• Top Highest Volatility. The sorted list of values based on the standard deviation of values for several alerts over time.<br>• Percentile. Objects are filtered based on the percentile entered.<br>Select the criteria for analysis of the metric that you select from the metric tree. |
| **Input Data** | |
| Objects | Select objects on which you want to base the widget data.<br>1. Click the **Add New Objects** icon and select objects in the pop-up window. The selected objects appear in a list in this section.<br><br>While selecting objects, you can use the **Filter** text box to search for objects. You can also expand the **Tag Filter** pane on the left hand side to select one or more object tag values. A list of objects with the selected tag values appears. If you select more than one value for the same tag, you can choose objects that have any of the tags applied. If you select more than one value for different tags, you can choose only the objects that have all the tags applied.<br>2. Optionally, select objects from the list and click the **Remove Selected Objects** icon to remove the selected objects.<br><br>Click the **Select All** icon to select all the objects in the list.<br><br>Click the **Clear Selection** icon to clear your selection of objects in the list. |
| All | If you select this option, the widget data is based on all the objects in your environment. The following sections provide options to refine the objects for the widget data. |
| **Input Transformation** | |
| Relationship | Transform the input for the widget based on the relationship of the objects. For example, if you select |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | the **Children** check box and a **Depth** of **1**, the child objects are the transformed inputs for the widget. |
| **Output Data** | |
| | Select an object type in your environment on which you want to base the widget data.<br>1. Click the **Add Object Type** icon to search for and add an object type.<br><br>When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box.<br>2. Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type.<br><br>If the objects have an input transformation applied, the transformed objects are the basis for the widget data. |
| Metric | Select a common metric or a metric for the selected object type in the list. The metric is the basis for the widget data. |
| Label | Type in a name that displays as a label for the metric.<br>You can add a label if you have selected **Metric Analysis › Top Highest Utilization** or **Metric Analysis › Top Lowest Utilization** as Top-N options in the **Configuration** section. |
| Unit | You can define measurement units for the metrics. Select a measurement unit in the **Unit** drop-down menu.<br>You can add a unit if you have selected **Metric Analysis › Top Highest Utilization** or **Metric Analysis › Top Lowest Utilization** as Top-N options in the **Configuration** section. |
| Maximum | Specify the maximum value based on which the bar size is calculated.<br>You can add a maximum value if you have selected any of the options under **Metric Analysis**. |
| Color Method | You can use the **Color Method** option to define a coloring criteria for each metric. If this option is set to **Custom**, you can enter color values in the **Yellow**, **Orange**, and **Red** text boxes. If you do not want to use color, select None.<br>You can add color thresholds if you have selected **Metric Analysis › Top Highest Utilization**, **Metric Analysis › Top Lowest Utilization**, or **Metric Analysis › Percentile** as Top-N options in the **Configuration** section. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | If the objects have an input transformation applied, you select tag values for the transformed objects. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>If the objects have an input transformation applied, you define filter criteria for the object types of the transformed objects.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** object type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |
| **Additional Columns** | |
| | Add metrics based on object types. The selected metrics are displayed as additional columns in the widget.<br><br>1. Click the **Add New Metrics** icon to add metrics based on object types. The metrics that you add appear in a list in this section.<br><br>While selecting object types for which you want to pick metrics, you can filter the object types by adapter type to pick an object type. On the metrics pane, click the **Select Object** icon to select an object for the object type. Pick metrics of the selected object from the metric tree.<br><br>For example, you can select the **Datacenter** object type, click the **Select Object** icon to display the list of data centers in your environment, and pick metrics of the selected data center.<br>2. Optionally, you can double-click a metric box in the list to customize the label of the metric and click **Update**. |

## Topology Graph Widget

The Topology Graph widget gives a graphical presentation of objects and their relationships in the inventory. You can customize each instance of the widget in your dashboard.

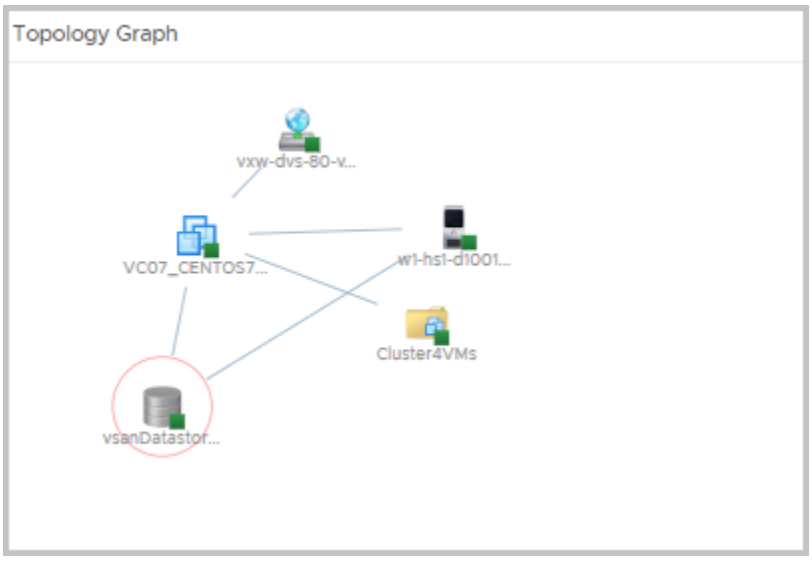


### How the Topology Graph Widget and Configuration Options Work

The Topology Graph widget helps you explore all nodes and paths connected to an object from your inventory. Connection between the objects might be a logical, physical, or network connection. The widget can display a graph that shows all of the nodes in the path between two objects, or that shows the objects related to a node in your inventory. You select the type of graph in the Exploration Mode when you configure the widget. You can select the levels of exploration between nodes in the displayed graph by using **Relationship** check boxes when you edit the widget. The widget displays all object types in the inventory by default, but you can select object types to view by using the Object View list during the configuration process. Double-clicking an object on the graph takes you to a detailed page about the object.

### Where You Find the Topology Graph Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Topology Graph Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

| Option | Description |
|---|---|
| Action | Use to select from predefined actions for each object type. To see available predefined actions, select an object in the graph and click the toolbar to select an action. For example, |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | when you select a datastore object in the graph, you can click **Delete Unused Snapshots for Datastore** to apply this action to the object. |
| Dashboard Navigation | Takes you to a predefined object. For example, when you select a datastore from the graph and click **Dashboard Navigation**, you can open the datastore in the vSphere Web Client. |
| Pan | Use to move the entire graph. |
| Show values on point | Provides a tool tip with parameters when you point to an object in the graph. |
| Zoom in | Zooms in the graph. |
| Zoom out | Zooms out the graph. |
| Hierarchical View | Use to switch to hierarchical view. Hierarchical view is activated only for Node Exploration mode and with selected inventory tree. |
| Graph View | Use to switch to graph view. |
| Object Detail | Select an object and click this icon to show the Object Detail page for the object. |
| Expand Node | Selects which object types related to your object to show on the graph. For example, if you select a virtual machine from the graph and click **Expand Node** toolbar icon and select **Host System**, the host on which the virtual machine is located is added to the graph. |
| Hide Node(s) | Use to remove a given object from the graph |
| Reset To Initial Object | Use to return to the initially displayed graph and configured object types. |
| Explore Node | Use to explore a node from a selected object in the graph. For example, if the graph displays a connection between a VM, a host, and a datastore, and you want to check the connection of the host with the other objects in the inventory, you can select the host and click **Explore Node**. |
| Status | Use to select objects based on their status or their state. |

**Topology Graph Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| Exploration Mode | Use **Node Exploration mode** to observe a selected object from an object list and the objects related to it. For example, if you select a virtual machine and select node exploration mode, the widget shows the host where the VM is placed and the datastore storing the files of the VM.<br><br>Use **Path Exploration mode** to observe the relation between two objects. You must select them from the Select First Object list and the Select Second Object list. For example, if you select to explore the path between a VM and a vCenter, the graph shows you both objects and all nodes in the path between the VM and server as datastore, datastore cluster, and data center.<br><br>**IMPORTANT**<br>To select object view is mandatory for the widget to start working in path exploration mode. |
| Show Paths | Use **All** to observe connections between a node and nodes related to it as well as connections between the nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph shows a VM connected to its datastore and host and the connection between the host and datastore. |
| | Use **Discovered Only** to observe directly related nodes. For example, if you are using node exploration mode and you select to observe a VM and all objects types, the graph will shows the VM connected to its datastore and to its host, but without the connection between the host and datastore. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Configuration File | The default configuration includes parent and child relationship. Drop-down options depend on the installed Solutions. You can add a new type of relationship to the Relationship pane. |
| Metric Configuration | Specifies a list with attributes to display. |
| Layout | Select whether you want a graph view or hierarchical view for the topology graph. |
| Tree type | For a hierarchical layout, select whether you want a tree type view. |
| **Input Data** | |
| Selected object | From the object list, select an object on which you want to base the widget data. |
| Degree of separation | Available only when node exploration mode is selected. Use to define the levels of exploration in node exploration mode. The lowest degree configuration shows only directly related nodes rather than higher degrees that show the inventory in details. |
| Select First Object | Available only in path exploration mode. Select the first object from the object list. |
| Select Second Object | Available only in path exploration mode. Select the second object from the object list. |
| Object view | Use to select which types of objects to observe in the graph. |
| Relationship | Select the type of relationship between objects to observe in the graph, respectively the details about your inventory . The common relationships for all objects are parent and child, but the list of relationships can vary depending on added solutions to VMware Aria OperationsVMware Cloud Foundation Operations. |

**View Widget**

The View widget provides the VMware Aria OperationsVMware Cloud Foundation Operations view functionality into your dashboard.

**How the View Widget and Configuration Options Work**

A view presents collected information for an object in a certain way depending on the view type. Each type of view helps you to interpret metrics, supermetrics, properties, alerts, policies, and data from a different perspective.

You can add the View widget to one or more custom dashboards and configure it to display data that is important to the dashboard users. List views can send interactions to other widgets.

**Where You Find the View Widget**

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the

dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

You can export the view as a CSV file for any view type.

**View Widget Toolbar Options**

The View widget toolbar depends on the displayed view type.

| Option | Description |
|---|---|
| Export as CSV | You can export the view as a CSV file for any view type. |
| Open in External Application | Ability to link to another application for information about the object. For example, you have a List view with VMs. You can select any VM and select **Open in External Application** to open the VM in vSphere Web Client. |
| Time Settings | Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.<br>• Relative Date Range. Select a relative date range of data transformation.<br>• Specific Date Range. Select a specific date range of data transformation.<br>• Absolute Date Range. Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br>The units of time available are: Hours, Days, Weeks, Months, and Years.<br><br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br>• Dashboard Time. Select this option to activate the dashboard time panel. The option chosen in the dashboard time panel is effective. The default time is 6 hours. |
| Near Real-Time Monitoring | Displays near real-time data collected at an interval of 20 seconds. Near real-time data is available between the 24 hour time range to upto three days. |
| Items per page | You can set the number of results that appear in the widget. Available for **List** view only. |
| Roll up interval | The time interval at which the data is rolled up. |
| Actions | An action on the selected object. Depends on the object type. |
| Filter | Limits the list to objects for a specific host, data center, and so on. You can drill-down in the hierarchical level. Available for **List**, **Trend**, and **Distribution** types of Views. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Filter by name | Limits the list to objects of a specific name. Available for **List** view only. |

**View Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Inventory trees | Select an existing predefined traversal spec to pick an object for the widget data. |
| Object | In self-provider mode, click the **Add Object** icon to select an object from the object list. The object list is displayed based on the inventory tree selection. You can also search for the object in this text box. |
| **Output Data** | |
| | A list of defined views available for the selected object is displayed. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | You can create, edit, delete, clone, export, and import views directly from the View widget configuration options.<br><br>For more information, see Views. |
| Auto Select First Row | Determines whether to start with the first row of data for list type views. |
| Show | Select one or more of the following items to display in the widget:<br>• To display the list of legends in the widget, select **Legend**.<br>• To display the name of the labels in the widget, select **Labels**. |

### Weather Map Widget

The Weather Map widget provides a graphical display of the changing values of a single metric for multiple resources over time. The widget uses colored icons to represent each value of the metric. Each icon location represents the metric value for particular resources. The color of an icon changes to show changes in the value of the metric.

#### How the Weather Map Widget and Configuration Options Work

You can add the Weather Map widget to one or more custom dashboards and configure it to display data that is important to different dashboard users. The data that appears in the widget is based on the configured options for each widget instance.

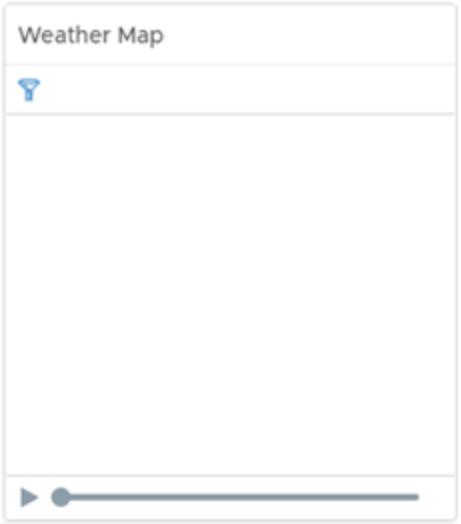


Watching how the map changes can help you understand how the performance of the metric varies over time for different resources. You can start or stop the display using the **Pause** and **Play** options at the bottom of the map. You can move the slider forwards or backwards to a specific frame in the map. If you leave the widget display and return, the slider remains in the same state.

The map does not show the real-time performance of the metrics. You select the time period, how fast the map refreshes, and the interval between readings. For example, you might have the widget play the metric values for the previous day, refreshing every half second, and have each change represent five minute's worth of metric values.

To view the object that an icon represents, click the object.

### Where You Find the Weather Map Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

### Weather Map Widget Toolbar Options

On the title bar of the widget, click the **Show Toolbar** icon to access the toolbar options.

The toolbar contains the icons that you can use to view the graph.

| Icon | Description |
|---|---|
| **Pause** and **Play** | Start or stop the display. The icon remains in the same state if you leave the widget display and return. |
| **Display Filtering Criteria** | View the current settings for the widget, including the current metric. |

### Weather Map Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Output Data** section provides options to select object types on which you are basing the widget data.

The **Output Filter** section provides options to restrict the widget data based on the selected filter criteria.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Redraw Rate | An interval at which cashed data is refreshed based on newly collected data.<br><br>For example, if you set metric history to `Last 6 hours` a nd image redraw rate to `15 minutes`, and data is collected every 5 minutes, the data collected during 10 minutes will not be calculated at the 15 minutes.<br><br>For example, if you set metric history to `Last 6 hours` a nd image redraw rate to `15 minutes`, and data is collected every 5 minutes, the data collected during 10 minutes will not be calculated at the 15 minutes. |
| Metric History | Select the time period for the weather map, from the previous hour to the last 30 days. |
| Metric Sample Increment | Select the interval between metric readings. For example, if you set this option to one minute and set the Metric History to one hour, the widget has a total of 60 readings for each metric. |
| Group by | Select a tag value by which to group the objects. |
| Sort by | Select **Object name** or **Metric value** to set the way to sort the objects. |
| Frame Transition Interval | Select how fast the icons change to show each new value. You can select the interval between frames and the number of frames per second (fps). |
| Start Over Delay | The number of seconds for the display to remain static when it reaches the end of the Metric History period, the most current readings, before it starts over again from the beginning. |
| Color | Shows the color range for high, intermediate, and low values. You can set each color and type minimum and maximum color values in the **Min Value** and **Max Value** text boxes.<br><br>If you leave the text boxes blank, VMware Aria OperationsV Mware Cloud Foundation Operations maps the highest and lowest values for the **Color By** metric to the end colors.<br><br>If you set a minimum or maximum value, any metric at or beyond that value appears in the end color. |
| **Output Data** | |
| | Select an object type in your environment on which you want to base the widget data.<br>1. Click the **Add Object Type** icon to search for and add an object type.<br><br>When you search for object types, you can filter the types in the list by selecting a type from the **Adapter Type** drop-down menu or by using the **Filter** text box. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 2. Optionally, select the object type from the list and click the **Delete Object Type** icon to remove the selected object type. |
| Metric | Select a common metric or a metric for the selected object type in the list. The metric will be the basis for the widget data. The object corresponding to the metric is the selected object for the widget. |
| **Output Filter** | |
| Basic | Pick tags to refine the widget data. The widget data is based on the objects that have the picked tags applied. If you pick more than one value for the same tag, the widget includes objects that have any of the tags applied. If you pick more than one value for different tags, the widget includes only the objects that have all the tags applied. |
| Advanced | Refine the widget data further based on the filter criteria for object types. The widget data is based on the objects for the filtered object types.<br><br>If the objects have a tag filter applied in the **Basic** subsecti on, you define filter criteria for the object types of the objects with tag filter applied. If the objects with tag filter applied do not belong to any of the object types in this filter criteria, the widget skips this filter and includes all the objects with tag filter applied.<br><br>1. In the first drop-down menu, select an object type.<br>2. In the second drop-down menu, select the option based on which you want to define the filter criteria. For example, if you select **Metrics** for the **Datacenter** objec t type, you can define a filter criteria based on the value of a specific metric for data centers.<br>3. In the drop-down menus and text boxes that appear, select or enter values to filter the objects.<br>4. To add more filter criteria, click **Add**.<br>5. To add another filter criteria set, click **Add another criteria set**. |

### Workload Widget

The Workload widget displays data indicating how hard a selected resource is working.

The Workload widget displays a graph depicting how hard the object that you selected is working. The Workload widget reports data on CPU usage, Memory usage, Disk I/O, and Network I/O.

#### Where You Find the Workload Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

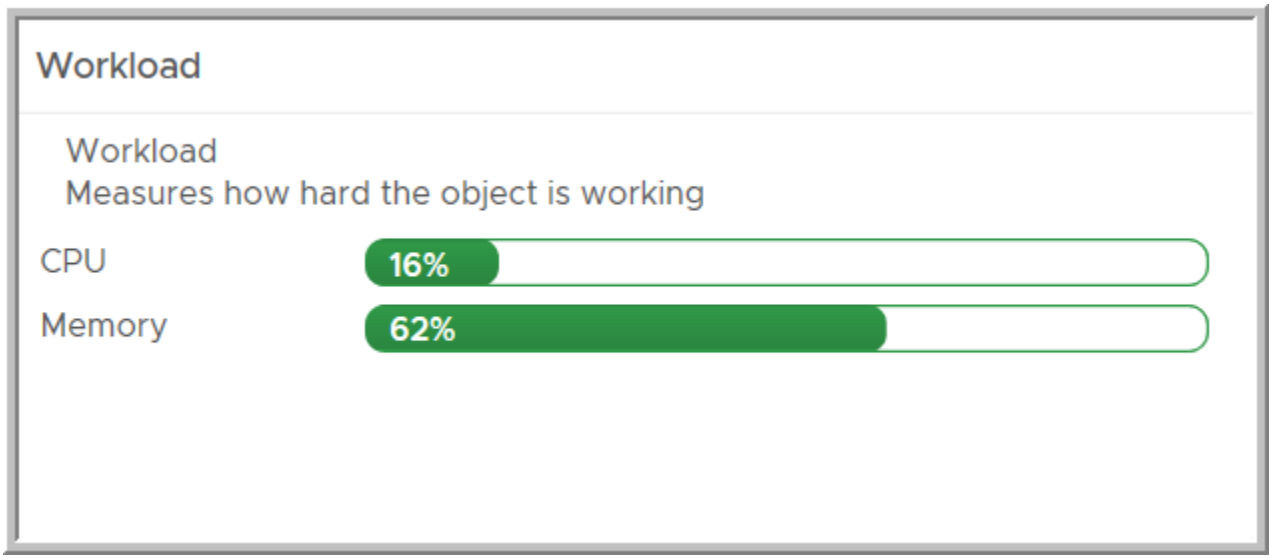


## About Datastore Metrics for Virtual SAN

The metric named `datastore|oio|workload` is not supported on Virtual SAN datastores. This metric depends on `datastore|demand_oio`, which is supported for Virtual SAN datastores.

The metric named `datastore|demand_oio` also depends on several other metrics for Virtual SAN datastores, one of which is not supported.

- The metrics named `devices|numberReadAveraged_average` and `devices|numberWriteAveraged_average` are supported.
- The metric named `devices|totalLatency_average` is not supported.

As a result, VMware Aria OperationsVMware Cloud Foundation Operations does not collect the metric named `datastore|oio|workload` for Virtual SAN datastores.

## Workload Widget Configuration Options

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

### Workload Pattern Widget

The Workload Pattern widget displays a historical view of the hourly workload of an object.

#### Where You Find the Workload Pattern Widget

The widget might be included on any of your custom dashboards. From the left menu, click **Operations › Dashboards** to see your configured dashboards.

To customize the data that appears in the dashboard widget, from the left menu, click **Operations › Dashboards**. To create your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, click **Create**. To edit your dashboard, from the left menu, click **Operations › Dashboards**. From the **Dashboards** panel, select the dashboard you want to edit and select **Actions › Edit**. Toggle between the **Views** and **Widgets** option to view and add a widget or view to the dashboard. The widgets list panel displays a list of all the predefined widgets. Drag a widget to the dashboard workspace in the upper panel.

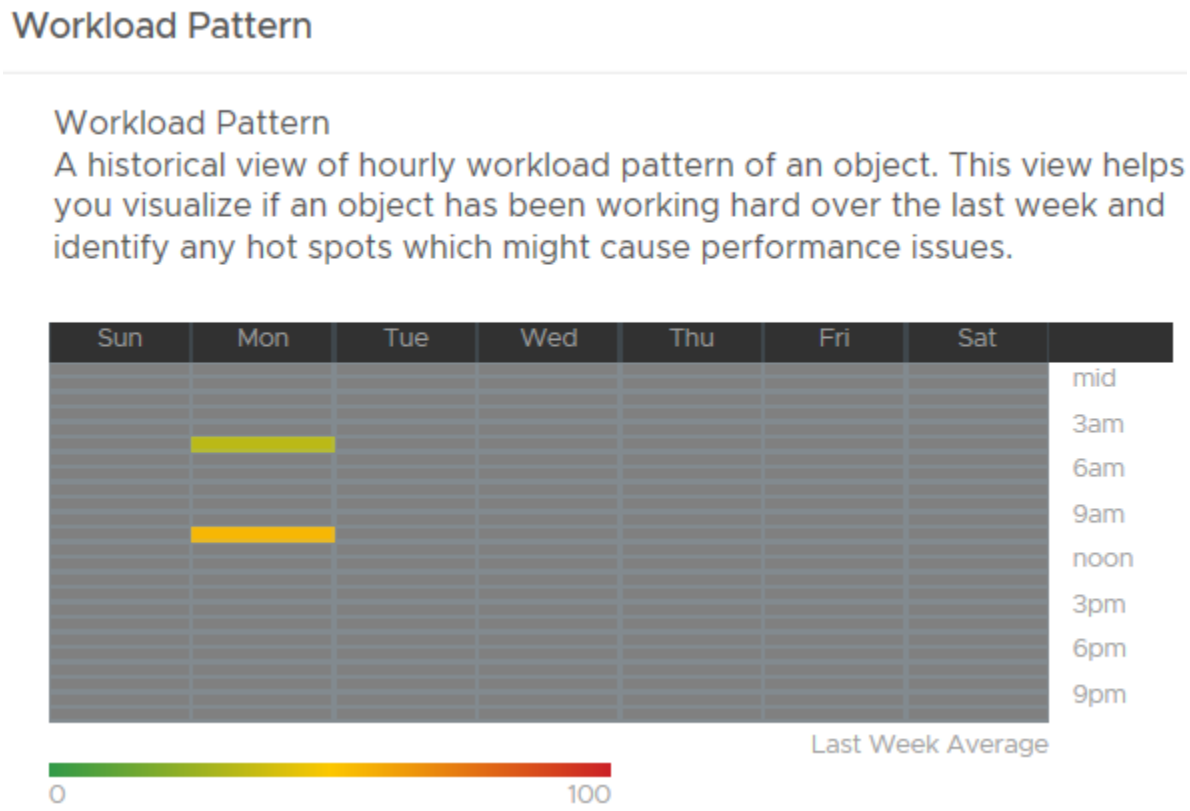


**Workload Pattern Widget Configuration Options**

On the title bar of the widget, click the **Edit Widget** icon to configure the widget.

The configuration options are grouped into one or more sections. You can select the objects on which you want to base the widget data and refine the objects in the following sections. Each section filters the objects further and pushes the filtered objects to the next section. The widget data is based on the objects that are the output of the last section.

The **Configuration** section provides general configuration options for the widget.

The **Input Data** section provides options to specify input for the widget. This section appears when the widget is in self provider mode.

| Option | Description |
|---|---|
| Title | Enter a custom title that identifies this widget from other instances that are based on the same widget template. |
| **Configuration** | |
| Refresh Content | Activate or deactivate the automatic refreshing of the data in this widget.<br><br>If not activated, the widget is updated only when the dashboard is opened or when you click the **Refresh** button on the widget in the dashboard. |
| Refresh Interval | If you activate the **Refresh Content** option, specify how often to refresh the data in this widget. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Self Provider | Indicates whether the objects for which data appears in the widget are defined in the widget or provided by another widget.<br>• On. You define the objects for which data appears in the widget.<br>• Off. You configure other widgets to provide the objects to the widget using the dashboard widget interactions options. |
| **Input Data** | |
| Object | Search for objects in your environment and select the object on which you are basing the widget data. You can also click the **Add Object** icon and select an object from the object list. You can use the **Filter** text box to refine the object list and the **Tag Filter** pane to select an object based on tag values. |

## Reports in VMware Aria OperationsVMware Cloud Foundation Operations

Configuring Reports
A report is a scheduled snapshot of views and dashboards. You can create reports in VMware Aria OperationsVMware Cloud Foundation Operations to represent objects and metrics. The report can contain a table of contents, cover page, and footer.

With the VMware Aria OperationsVMware Cloud Foundation Operations reporting functions, you can generate a report to capture details related to current or predicted resource needs. You can download the report in a PDF or CSV file format for future and offline needs.

## Types of Report Templates

There are two types of report templates in VMware Aria OperationsVMware Cloud Foundation Operations.
- Predefined report templates that are out of the box with VMware Aria OperationsVMware Cloud Foundation Operations. For more information, see Accessing Report Templates .
- Custom report templates that you create based on your requirement. For more information, see Create a Report Template.

## Create a Report Template

You create a report to generate a scheduled snapshot of views and dashboards. You can track current resources and predict potential risks to the environment. You can schedule automated reports at regular intervals.The name and description of the report template as they appear in the list of templates on the **Report Templates** tab.

1. To create report templates, from the left menu, click **Operations › Reports**.
2. From the **Reports** panel, click **Create**. From the **Create Report Template** page, complete the options in each tab.
3. At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the report template.
4. After you have added all the details, click **Create** to create the report template.

**Name and Description Tab**

**Table 207: Name and Description Options in the Create Report Template Page**

| Option | Description |
|---|---|
| Name | Name of the template as it appears on the **Report Templates** tab. |
| Description | Description of the template. |

**Report Content Tab**

The report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. You can combine different views and dashboards and order them to suit your needs. The report template contains images of the views/dashboards that are added.

**Table 208: Views and Dashboards Options and Sections in the Create Report Template Page**

| Option | Description |
|---|---|
| Report Template Structure | To add a view or a dashboard to your report template, select it from the **Views and Dashboards** list in the right pane and drag it to the **Report Template Structure** pane. |
| View and Dashboards | Select **Views** or **Dashboards** to display a list of available views or dashboards that you can add to the template. |
| Filter | Search for views or dashboards by name. To see the complete list of views or dashboards, delete the search box contents and press Enter. |
| Vertical Ellipsis > Portrait/Landscape | You can select a portrait or landscape orientation for each view or dashboard from the vertical ellipsis next to the title of the view/dashboard after you drag and drop the view/ dashboard to the left pane. |
| Vertical Ellipsis > Colorization | You can activate or deactivate a colorized PDF output for each list view from the vertical ellipsis next to the title of a list view after you drag and drop the list view to the left pane.<br>Available only for list views. |

**Layout and Format Tab**

The report template can contain layout options such as a cover page, table of contents, and footer. Formats are the outputs in which you can generate the report.

**Table 209: Layout and Format Options in the Create Report Template Page**

| Option | Description |
|---|---|
| Cover Page | Can contain an image up to 5 MB.<br><br>The default report size is 8.5 inches by 11 inches. The image is resized to fit the report front page. |
| Table of contents | Provides a list of the template parts, organized in the order of their appearance in the report. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Footer | Includes the date when the report is created, a note that the report is created by VMware Aria OperationsVMware Cloud Foundation Operations, and page number. |
| PDF | With the PDF format, you can read the reports, either on or off line. This format provides a page-by-page view of the reports, as they appear in printed form. |
| CSV | In the CSV format, the data is in a structured table of lists. |

## Manage Report Templates

Using VMware Aria Operations you can manage report templates by running, scheduling, and generating report templates. Additionally, you can create custom report templates, edit existing report templates, download report templates, and so on.

**Reports Workflow**

The following flowchart describes a simple workflow for reports in VMware Aria OperationsVMware Cloud Foundation Operations.

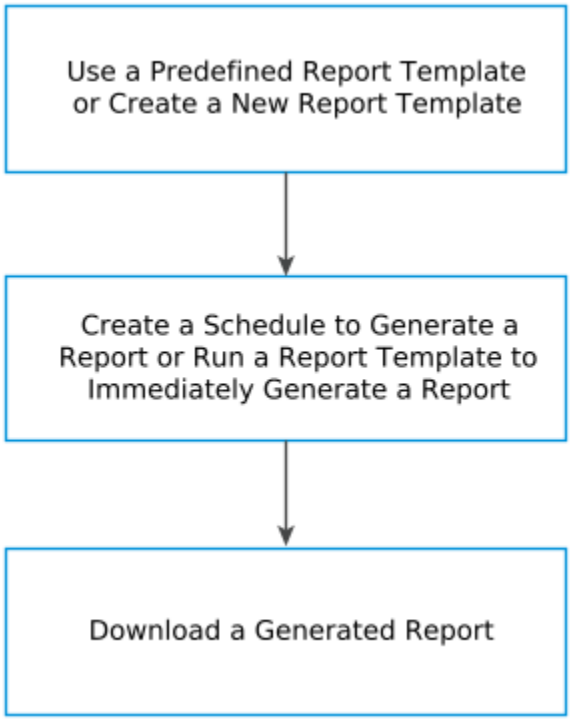


### Accessing Report Templates

A report template contains views and dashboards. Views present collected information for an object. Dashboards give a visual overview of the performance and state of objects in your virtual infrastructure. VMware Aria OperationsVMware Cloud Foundation Operations offers several predefined report templates that you can use based on your requirement.

**Where You Can Access Report Templates From**

From the left menu, click **Operations › Reports**. The **Report Templates** page is in the right panel, or

The listed report templates are user-defined and predefined by VMware Aria OperationsVMware Cloud Foundation Operations. You can order them by template name, description, subject, date they were modified, last run report, or the user who modified them. For each template, you can see the number of generated reports and schedules.

You can filter the reports based on the name of the report template, the subject, and the owner. You can click **Add** to create a report template. For information about creating a report template, see Create a Report Template.

You can select a report template from the list, click the vertical ellipsis against each report template, and select options such as run, edit, schedule, delete, clone, and export a report.

**Table 210: Predefined Filter Groups**

| Filter Group | Description |
|---|---|
| Name | Filter by the template name. For example, type `my template` to list all reports that contain the `my template` phrase in their name. |
| Subject | Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by the other objects. |
| Owner | Filter by the owner of the report template. |

The maximum number of reports per template is 10. After the tenth report is generated, VMware Aria OperationsVMware Cloud Foundation Operations deletes the oldest report. vSphere users must be logged in until the report generation is complete. If you log out or your session expires, the report generation fails.

**Report Template Actions**

You can select more than one report template and perform a set of actions by clicking the horizontal ellipsis next to the **Add** option.

| Option | Description |
|---|---|
| Delete | Deletes the report template. |
| Export | Downloads the report template. |
| Import | Allows you to import a report template by selecting a report template in XML or zip file format.<br>To import a report template:<br>• Click the **Import** option from the horizontal ellipsis.<br>• Click **Browse** and select a report template ZIP or XML file to import.<br>• Select if you want to Overwrite or Skip the file in case of a conflict.<br>• Click **Import** to import the report template, and click **Done**. |
| Change default cover image | Allows you to change the default cover image of the report template. For more information, see Upload a Default Cover Page Image for Reports. |

### Schedule a Report

In VMware Aria Operations, the schedule of a report is the time and recurrence of report generation. To generate a report on a selected date, time, and recurrence, you create a schedule for the report template. You set the email options to send the generated report to your team.

- Download the generated report to verify the output.
- To activate sending email reports, you must have configured Outbound Alert Settings. See Notifications in .

The date range for the generated report is based on the time when the report is generated and not on the time when you schedule the report or when VMware Aria OperationsVMware Cloud Foundation Operations places the report in the queue.

> **NOTE**
> Only users created in VMware Aria OperationsVMware Cloud Foundation Operations can add and edit report schedules.

1. From the left menu, click **Operations › Reports**.
2. From the Report Template page, select the relevant report template from the list.
3. Click the vertical ellipsis and select **Schedule**.
4. Select an object and click **Next**.
5. Select the time zone, date, hour, and minutes (in the range of 0, 15, 30, and 45 minutes) to start the report generation.

   VMware Aria OperationsVMware Cloud Foundation Operations generates the scheduled reports in sequential order. Generating a report can take several hours. This process might delay the start time of a report when the previous report takes an extended period of time.

6. From the **Recurrence** drop-down menu, select one of the following options for report generation:

| Option | Description |
|---|---|
| Daily | You can set the periodicity in days. For example, you can set report generation to every two days. |
| Weekly | You can set the periodicity in weeks. For example, you can set report generation to every two weeks on Monday. |
| Monthly | You can set the periodicity in months. |

7. Select the **Email report** check box to send an email with the generated report.

   Email a generated report to a predefined email group or to a network shared location. For more information about how to set up and configure the email options, see Add a Standard Email Plug-In for Outbound Alerts.

   a) In the **Email addresses** text box, enter the email addresses that must receive the report. You can also add email addresses in the CC list and BCC list.
   b) Select an outbound rule.

   An email is sent according to this schedule every time a report is generated.

8. Save a generated report to an external location.

   For more information about how to configure an external location, see Add a Network Share Plug-In for Reports

9. You can add a relative path to upload the report to a predefined sub folder of the Network Share Root folder. For example, to upload the report to the share host `C:/documents/uploadedReports/SubFolder1`, in the **Relative Path** text box, enter `SubFolder1`. To upload the report to the Network Share Root folder, leave the **Relative Path** text box empty.

### Editing a Report Schedule

To edit the schedule of a report, click the link in the **Schedules** column against the report template from the **Report Templates** page, and then from the **Scheduled Reports** dialog box, click **Edit Schedule**. You see the **Scheduled Reports** page.

**Table 211: Scheduled Reports Toolbar Options**

| Options | Description |
|---|---|
| New Schedule | You can create a schedule for the report. |
| Edit Schedule | You can edit an existing report schedule. |
| Delete Schedule | You can delete an existing report schedule. |
| Transport Report Schedule | You can assign a new owner for the selected report schedule. You can select a target user from the **Transfer Report Schedules** dialog box. |

**NOTE**
You can edit, clone, and delete report templates. Before you do, familiarize yourself with the consequences of these actions. When you edit a report template and delete it, all reports generated from the original and the edited templates are deleted. When you clone a report template, the changes that you make to the clone do not affect the source template. When you delete a report template, all generated reports are also deleted.

### Generate and Regenerate a Report

To generate a report, use a predefined or custom report template.

Create a report template.

1. From the left menu, click **Operations › Reports** .
2. From the **Reports** panel, select **Manage**.
3. From the **Report Templates** page on the right, navigate to the relevant report template, click the vertical ellipsis, and select **Run**.
4. Select an associated object from the **Select Object** dialog box and click **OK**.

The report is generated and listed on the **Generated Reports** tab.

**NOTE**
To regenerate the selected report, from the **Generated Reports** tab, click the vertical ellipsis against the generated report and select **Run**.

Download the generated report and verify the output.

### Accessing Generated Reports

You can view a list of report templates that have been generated in VMware Aria OperationsVMware Cloud Foundation Operations.

#### Where You Can Access Generated Reports From

From the left menu, click **Operations › Reports**. From the **Reports** panel, click **Generated Reports**. The right pane contains all the generated reports, or

If the report is generated through a schedule, the owner is the user who created the schedule.

**NOTE**
The maximum number of reports per template is 10. After the tenth report is generated, VMware Aria OperationsVMware Cloud Foundation Operations deletes the oldest report.

To select a generated report from the list, click the vertical ellipsis against each generated report and select options such as run and delete. You can also select more than one generated report and select the **Delete** button above the data grid to delete a generated report.

You can filter the reports list by adding a filter from the upper-right corner of the panel.

**Table 212: Predefined Filter Groups**

| Filter Group | Description |
|---|---|
| Report Name | Filter by the report template name. For example, type `my template` to list all reports that contain the `my template` phrase in their name. |
| Template | Filter by the report template. You can select a template from a list of templates applicable for this object. |
| Completion Date/Time | Filter by the date, time, or time range. |
| Subject | Filter by another object. If the report contains more than one view applicable for another type of object, you can filter by that second object. |
| Status | Filter by the status of the report. |

You can download a report in a PDF or CSV format. You define the format that a report is generated in the report template.

If you log in to VMware Aria OperationsVMware Cloud Foundation Operations with vCenter credentials and generate a report, the generated report is always blank.

**Download a Report**

To verify that the information appears as expected, you download the generated report.

Generate a report.

1. From the left menu, select **Operations › Reports** .
2. From the **Reports** panel, click **Generated Reports**.
3. From the **Generated Reports** page on the right, click the PDF or the CSV icon in the **Download** column to download the report.

VMware Aria OperationsVMware Cloud Foundation Operations saves the report file.

Schedule a report generation and set the email options, so your team receives the report.

### Upload a Default Cover Page Image for Reports

You can upload a common default image for the cover page of reports. You do not have to upload a cover page for each report. The cover pages of predefined reports are modified when you use this option. The cover pages of user-defined reports do not change.

**Where Do You Upload a Default Cover Page Image for Reports**

To upload a default cover page for reports, from the left menu, click **Operations › Report**. From the **Reports** panel, click **Manage**. From the **Report Templates** page on the right side, click the horizontal ellipsis next to the **Add** option and click the **Change default cover image** option.

**How Do You Upload a Default Cover Page Image for Reports**

Browse for the image that you want to add to the cover page and click **Save**. You can also use the default product image that is available.

### Add a Network Share Plug-In for VMware Aria OperationsVMware Cloud Foundation Operations Reports

You add a Network Share plug-in when you want to configure VMware Aria OperationsVMware Cloud Foundation Operations to send reports to a shared location. The Network Share plug-in supports only SMB version 2.1.

Verify that you have read, write, and delete permissions to the network share location.

1. From the left menu, click **Operations › Configurations**, and then click the **Outbound Settings** tile.
2. Click **Add**, and from the **Plug-In Type** drop-down menu, select **Network Share Plug-in**.
   The dialog box expands to include your plug-in instance settings.
3. Enter an **Instance Name**.

   This is the name that identifies this instance that you select when you later configure notification rules.
4. Configure the Network Share options appropriate for your environment.

| Option | Description |
|---|---|
| Domain | Your shared network domain address. |
| User Name | The domain user account that is used to connect to the network. |
| Password | The password for the domain user account. |
| Network share root | The path to the root folder where you want to save the reports. You can specify subfolders for each report when you configure the schedule publication.<br><br>You must enter an IP address. For example, `\\IP_address\ShareRoot`. You can use the host name instead of the IP address if the host name is resolved to an IPv4 when accessed from the VMware Aria OperationsVMware Cloud Foundation Operations host.<br><br>**NOTE**<br>Verify that the root destination folder exists. If the folder is missing, the Network Share plug-in logs an error after 5 unsuccessful attempts. |

5. Click **Test** to verify the specified paths, credentials, and permissions.

   The test might take up to a minute.
6. Click **Save**.
   The outbound service for this plug-in starts automatically.
7. To stop an outbound service, select an instance and click **Deactivate** on the toolbar.

This instance of the Network Share plug-in is configured and running.

Create a report schedule and configure it to send reports to your shared folder.

## Views in VMware Aria OperationsVMware Cloud Foundation Operations

Configuring Views
VMware Aria OperationsVMware Cloud Foundation Operations provides several types of views. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective. Views also show information that the adapters in your environment provide.

You can configure VMware Aria OperationsVMware Cloud Foundation Operations views to show transformation, forecast, and trend calculations.

- The transformation type determines how the values are aggregated.
- The trend option shows how the values tend to change, based on the historical, raw data. The trend calculations depend on the transformation type and roll up interval.
- The forecast option shows what the future values can be, based on the trend calculations of the historical data.

You can use VMware Aria OperationsVMware Cloud Foundation Operations views in different areas of VMware Aria OperationsVMware Cloud Foundation Operations.

- To manage all views, from the left menu, click **Operations › Views**. From the **Views** panel, click **Manage**.
- To see the data that a view provides for a specific object, navigate to that object, click the **Details** tab, and click **Views**.
- To see the data that a view provides in your dashboard, add the View widget to the dashboard. For more information, see View Widget.

**Table 213: Options from the Views panel**

| Options | Description |
|---|---|
| Manage | You can manage views by clicking **Operations › Views**. From the **Views** panel, click **Manage**. |
| Create | Use this option to create a view. See Create and Configure a View. |
| Search | You can search for a view across the **Recents** and **All** folders in the **Views** panel. |
| Recent | The views are listed in the order in which you select them, with the most recent view that you selected, appearing at the top. Up to ten views can be displayed as Recent views.<br>If you do not pin the view and log out of the user interface, on logging back in, the view is removed from the Recents folder |
| All | Lists the views based on their type. You can use this menu for quick navigation through your views. When you navigate to a view using the **Operations › Views** option, the views are listed in the **Views** panel under **All**. You can also search for views using keywords and letters. |

## Views and Reports Ownership

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

## Accessing Predefined Views

You can access some of the useful, predefined views from the **Views** home page.

To access these views, from the left menu, click **Operations › Views**. From the **Views** panel, click **Home**.

The views are categorized as follows: Availability, Capacity, Configuration, Inventory, Performance, and Compliance. To easily access some of the useful, predefined dashboards under these categories, click on the drop-down button against the selected category and click on the specific dashboard.

## Views Overview

A view presents collected information for an object in a certain way depending on the view type. Each type of view helps you to interpret metrics, properties, policies of various monitored objects including alerts, symptoms, and so on, from a different perspective.

### How You Access the Views Page

From the left menu, click **Operations › Views**. From the **Views** panel, click **Manage** to access the **Views** page.

### Manage and Preview Views

You can preview a view by clicking a view from the **Views** page. Add an object if necessary by clicking **Preview source** from the upper-right corner of the **Views** page. The preview of the view appears just below the **Views** option in the right pane.

You can select a view from the list, click the vertical ellipsis against each view, and select the various options such as edit, delete, clone, and export a view.

You can filter the views based on the name, type, description, subject, and owner. You can click the **Add** option to create a view. For information about creating a view, see Create and Configure a View.

Views are also categorized and listed in the **Views** panel based on the type of view and subject.

**Table 214: Filter Groups**

| Filter Group | Description |
|---|---|
| Name | Filter by the view name. For example, type `my view` to list all views that contain the `my view` phrase in their name. |
| Type | Filter by the view type. |

*Table continued on next page*

*Continued from previous page*

| Filter Group | Description |
|---|---|
| Description | Filter by the view description. For example, type `my view` to list all views that contain the `my view` phrase in their description. |
| Subject | Filter by the subject. |
| Owner | Filter by owner. |

**Datagrid Options**

| Column Names | Description |
|---|---|
| Name | Displays the name of the view. |
| Type | Displays the type of view: list, summary, trend, distribution, text, or image. |
| Description | Displays the description of the dashboard. |
| Subject | Displays the base object type for which the view shows information. |
| Dashboard Usage | Displays the number of dashboards where the view is used. Click the number in the column to view the name of the dashboard/s. Click the dashboard name to navigate to the dashboard. |
| Report Usage | Displays the number of reports where the view is used. Click the number in the column to view the name of the report/s. Click the report name to navigate to the report template in edit mode. |
| Last Modified | Displays the date the view was last modified. |
| Modified By | Displays the user who last modified the view. |

**Views Actions**

You can select more than one view and perform a set of actions by clicking the horizontal ellipsis next to the **Add** option.

| Option | Description |
|---|---|
| Delete | Deletes the view. |
| Export | Downloads the view. |
| Import | Allows you to import a view by selecting a view in XML or zip file format.<br>To import a view:<br>• Click the **Import** option from the horizontal ellipsis.<br>• Click **Browse** and select a view XML or ZIP file to import.<br>• Select if you want to Overwrite or Skip the file in case of a conflict.<br>• Click **Import** to import the view, and click **Done**. |

## Views and Reports Ownership

The owner of views, reports, or templates might change over time.

The default owner of all predefined views and templates is System. If you edit them, you become the owner. If you want to keep the original predefined view or template, you have to clone it. After you clone it, you become the owner of the clone.

The last user who edited a view, template, or schedule is the owner. For example, if you create a view you are listed as its owner. If another user edits your view, that user becomes the owner listed in the Owner column.

The user who imports the view or template is its owner, even if the view is initially created by someone else. For example, *User 1* creates a template and exports it. *User 2* imports it in back, the owner of the template becomes *User 2*.

The user who generated the report is its owner, regardless of who owns the template. If a report is generated from a schedule, the user who created the schedule is the owner of the generated report. For example, if *User 1* creates a template and *User 2* creates a schedule for this template, the generated report owner is *User 2*.

## Create and Configure a View

To collect and display information for a specific object, you can create a custom view.

1. From the left menu, click **Operations › Views**.
2. From the **Views** panel, click **Create**.
3. Select one of the following views from the right panel.
   - List View
   - Summary View
   - Trend View
   - Distribution View
   - Text View
   - Image View
4. At the end of each tab in the selected view, you can go to the previous or next tab. You can also cancel the creation of the view.
5. After you have added all the details, click **Create** to create the view.

### List View

List views provide tabular data about specific objects in the monitored environment that correspond to the selected view.

#### Where You Find the List View

From the left menu, click **Operations › Views**. From the **Views** panel, click **Create**. Click **List** from the right panel.

#### Name and Configuration Tab

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Items per page | Select the number of items per page. Each item is one row and its metrics and properties are the columns. |
| Top result count | Select the top results. Restricts the number of results. For example, if you list all the clusters in a View, selecting 10 in this option displays the top 10 clusters with the relevant |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | information. You can reduce the number of rows for the purposes of reporting. |
| Include Deleted Objects | Select to add deleted objects. |
| Show Objects | Select the type of object you want displayed in the view. You can select **Existing**, **Deleted**, or **All** objects. |
| Show Object Creation Date | Select to display the date the object was created. |
| Make the view available at > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available at > Report template creation and modification | Select if you want to make the view available in a report template. |
| Make the view available at > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br><br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

**Data Tab**

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria OperationsVMware Cloud Foundation Operations collects, calculates, and presents the information for the view.

**How to Add Data to a View**

If you selected more than one subject, click the subject for which you add data. Double-click either a metric or a property from the tree in the left panel to add it to the view. For each subject that you select, the data available to add might be different. The Data, Transformation, and Configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Preview Source**.

| Option | Description |
|---|---|
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| Group By | You can group the results based on a parent object, by making a selection in the Group By drop-down option. If you generate a report based on the list view for which a group has been specified, the report displays group-based information for the selected object. You can also view summary calculations for the group of objects in the report, along with the total summary results for all the objects. |
| **Breakdown By** | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Add interval breakdown | Select this check box to see the data for the selected resources broken down in time intervals.<br>After you select this check box, you can enter a label, specify whether the values have to be sorted in ascending or descending order, and select a breakdown interval for the time range. |
| Add Instance breakdown | Select this check box to see the data for all instances of the selected resources.<br>After you select this check box, you can enter a label and select a metric group to break down all the instances in that group. Deselect **Show non-instance aggregate metric** to display only the separate instances. Deselect **Show only instance name** to display the metric group name and instance name in the instance breakdown column.<br><br>For example, you can create a view to display CPU usage by selecting the metric **CPU:0\|Usage**. If you add an instance breakdown column, the column **CPU:0\|Usage** displays the usage of all CPU instances on separate rows (0, 1, and so on). To avoid ambiguity, you can change the metric label of **CPU:0\|Usage** to **Usage**. |
| **Data Grid options** | |
| **Self** drop down option | Click the drop down button to select an ancestor or descendant of the selected subject. You can select metrics and properties from the data selection tree and then configure them.<br><br>**NOTE**<br>Error messages are displayed in the following cases:<br>• If the same object type is used as an ancestor/descendant and as the primary object type.<br>• If you do not select a metric for the primary object type. You must select at least one metric.<br>• If the number of related ancestor/descendant object types used in the configuration exceeds a maximum of 3. |
| Add Object Name | If you have selected an ancestor or descendant from the **Self** drop down option, click **Add Object Name** to add the name of the selected descendant or ancestor as a column name in the **Preview Source** pane. |
| Data selection tree (Metrics and Properties) | Select a metric or a property. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>The metrics presented by default are a subset of available metrics. If the desired metric is not represented, use the **Select Object** button in the title bar of the data selection tree to re-filter the displayed list. |
| Data column | Click the metric or property to enter configuration details in the configuration column. |
| Transformation column | Displays the type of transformation applied to the data. |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand (MHz) from the **Units** drop-down menu, you can change the value to Hz, kHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |
| Sort order | Orders the values in ascending or descending order<br>**NOTE**<br>Sort order is not activated for ancestor or descendant objects. |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>• Minimum. The minimum value of the metric over the selected time range.<br>• Maximum. The maximum value of the metric over the selected time range.<br>• Average. The mean of all the metric values over the selected time range.<br>• Sum. The sum of the metric values over the selected time range.<br>• First. The first metric value for the selected time range.<br>• Last. The last value of a metric within the selected time range.<br>If you have selected **Last** as the transformation in versions before vRealize Operations 6.7, and the end of specified time range is not before the last five minutes, use the **Current** transformation.<br>• Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>• Standard Deviation. The standard deviation of the metric values.<br>• Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | displays the value for memory.usage when cpu.usage is at a maximum.<br>• Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>• Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>• Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, `sum/((max + min)/2)`. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps.<br>• Timestamp: You can choose between Absolute Timestamp or Relative Timestamp.<br>– If applied to a numeric metric/property defined with a time-unit definition, the actual value is converted to a human readable timestamp. The metric value is rounded-off to an hour.<br>– In the remaining cases, a timestamp is displayed when metrics and properties are added or modified. In this case, the behavior is the same as the Timestamp option selected for a non-Timestamp transformation. Applicable for Absolute Timestamp and Relative Timestamp.<br>Available for List view and Minimum, Maximum, Current, First, and Last transformation. |
| Ranges for metric coloring | You can associate colors to metrics by entering a percentage, range, or specific state. For example, you can enter Powered Off in the **Red Bound** field when you select virtual machine as an object. You can set the colors only for views and not for csv or pdf formats. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select Sum as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | This option is applicable to the Transformation configuration option. |

**Time Settings Tab**

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. Data is collected based on the browser time.

**Table 215: Time Settings Options**

| Configuration Option | Description |
|---|---|
| Time Range Mode | In Basic mode, you can select date ranges.<br><br>In Advanced mode, you can select any combination of relative or specific start and end dates.<br><br>You can also activate the **Business Hour** option and select business hours/days for weekdays. |
| Relative Date Range | Select a relative date range of data transformation.<br><br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br><br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br><br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br><br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br><br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br><br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br><br>Available in Advanced mode. |

*Table continued on next page*

*Continued from previous page*

| Configuration Option | Description |
|---|---|
| Specific Start Date | Select a specific start date of data transformation.<br><br>Available in Advanced mode. |
| Specific End Date | Select a specific end date of data transformation.<br><br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |
| Business Hours | Select business hours from Monday to Sunday by moving the sliders on the left and right sides to set the start and end time for each day of the week.<br>For example, as a VM owner, you can track the average utilization of VMs over a week (business days), during specified hours of the day (business hours).<br><br>This option is available for Minimum, Maximum, Average, Sum, and Percentile transformations<br><br>Available in Advanced mode for List Views. |

**Filter Tab**

The filter option allows you to add additional criteria when the view displays too much information. For example, a List view shows information about the health of virtual machines. From the **Filter** tab, you add a risk metric less than 50%. The view displays the health of all virtual machines with risk less than 50%. For selected criteria you can also apply Business Hours, if the selected transformation type you add as a filter is supported by the business hours functionality.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**. You can activate Business Hours for the metric selected.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

> **NOTE**
> For Symptom based views, in the filter, if you select either **Symptom Definition Name**, **Alert Definition Name**, or **Alert Type**, it is recommended that you select a preview source that has a smaller number of objects.

**Table 216: Filter Add Options**

| Option | Description |
|---|---|
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric or property, all the instances of the object for which the criteria is met, will be displayed in the preview screen. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | For instance metrics or properties, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, Sum, and Timestamp. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

**Summary Tab**

You can add more than one summary row or column and configure each to show different aggregations. In the summary configuration panel, you select the aggregation method and what data to include or exclude from the calculations.

To add a summary row or column to a view, from an existing or new view dialog box, click the **Summary** tab in the right pane. Click the plus sign to add a summary row.

For the Summary view, the summary column shows aggregated information by the items provided on the **Data** tab.

**Previous, Next, Create, and Cancel Options**

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

**Summary View**

Summary views provide tabular data about the use of resources in the monitored environment.

**Where You Find the Summary View**

From the left menu, click **Operations › Views**. From the **Views** panel, click **Create**. Click **Summary** from the right panel.

**Name and Configuration Tab**

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Items per page | Select the number of items per page. Each item is one row and its metrics and properties are the columns. |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br><br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | objects. You add a data center from the list. The view is not visible anymore on the data center level. |

**Data Tab**

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria OperationsVMware Cloud Foundation Operations collects, calculates, and presents the information for the view.

**How to Add Data to a View**

If you selected more than one subject, click on the subject for which you want to add data. Double-click either a metric or a property from the tree in the left panel to add it to the view. For each subject that you select, the data available to add might be different. The Data, Transformation, and Configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
|---|---|
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |
| Data selection tree | Select a metric or property |
| Data column | Click the metric or property to enter configuration details in the configuration column. |
| Transformation column | |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\|Demand(MHz) from the **Units** drop-down menu, you can change the value to Hz, KHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |
| Sort order | Orders the values in ascending or descending order. |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>• Minimum. The minimum value of the metric over the selected time range.<br>• Maximum. The maximum value of the metric over the selected time range.<br>• Average. The mean of all the metric values over the selected time range.<br>• Sum. The sum of the metric values over the selected time range. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • First. The first metric value for the selected time range.<br>• Last. The last value of a metric within the selected time range.<br>If you have selected **Last** as the transformation in versions before vRealize Operations 6.7, and the end of specified time range is not before the last five minutes, use the **Current** transformation.<br>• Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>• Standard Deviation. The standard deviation of the metric values.<br>• Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum.<br>• Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>• Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>• Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, `sum/((max + min)/2)`. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select Sum as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.<br><br>This option is applicable to the Transformation configuration option. |

**Time Settings Tab**

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. Data is calculated based on the browser time.

**Table 217: Time Settings Options**

| Configuration Option | Description |
|---|---|
| Time Range Mode | In Basic mode, you can select date ranges.<br><br>In Advanced mode, you can select any combination of relative or specific start and end dates.<br><br>You can also activate the **Business Hour** option and select business hours/days for weekdays. |
| Relative Date Range | Select a relative date range of data transformation.<br><br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br><br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br><br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br><br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br><br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br><br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br><br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br><br>Available in Advanced mode. |
| Specific End Date | Select a specific end date of data transformation.<br><br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

### Filter Tab

The filter option allows you to add additional criteria when the view displays too much information. For example, a view shows information about the health of virtual machines. From the **Filter** tab, you add a risk metric less than 50%. The view displays the health of all virtual machines with risk less than 50%.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

**Table 218: Filter Add Options**

| Option | Description |
|---|---|
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric or property, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br><br>For instance metrics or properties, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, Sum, and Timestamp. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

### Summary Tab

You can add more than one summary row or column and configure each to show different aggregations. In the summary configuration panel, you select the aggregation method and what data to include or exclude from the calculations.

To add a summary row or column to a view, from an existing or new view dialog box, click the **Summary** tab in the right pane. Click the plus sign to add a summary row.

For the Summary view, the summary column shows aggregated information by the items provided on the **Data** tab.

### Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

## Trend View

Trend views use historic data to generate trends and forecasts for resource use and availability in the monitored environment.

### Where You Find the Trend View

From the left menu, click **Operations › Views**. From the **Views** panel, click **Create**. Click **Trend** from the right panel.

**Name and Configuration Tab**

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| The maximum plot lines | Enter the maximum number of plot lines. Limits the output in terms of the objects displayed in the live preview of the view type on the left upper pane. The number you set as the maximum number of plot lines determines the plot lines.<br><br>For example, if you plot historical data and set the maximum at 30 plot lines, then 30 objects are displayed. If you plot historical, trend, and forecast lines, and set the maximum to 30 plot lines, then only 10 objects are displayed as each object has three plot lines. |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br><br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

**Data Tab**

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria OperationsVMware Cloud Foundation Operations collects, calculates, and presents the information for the view.

**How to Add Data to a View**

If you selected more than one subject, click on the subject for which you want to add data. Double-click the data from the tree in the left panel to add it to the view. For each subject the data available to add, might be different.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
|---|---|
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Data selection tree | Select a metric or a property. |
| Data column | Click the metric or property to enter configuration details in the configuration column. |
| Transformation column | Displays the type of transformation applied to the data. |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand (MHz) from the Units drop-down menu, you can change the value to Hz, KHz, or GHz. If you select Auto, the scaling is set to a meaningful unit. |
| Data Series | You can select whether to include historical data, trend of historical data, and forecast for future time in the trend view calculations. |
| Ranges for metric coloring | You can associate colors to metrics by entering a percentage, range, or specific state. For example, you can enter Powered Off in the **Red Bound** field when you select virtual machine as an object. You can set the colors only for views and not for csv or pdf formats. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select<br>`Sum`<br>as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.<br><br>This option is applicable to the Transformation configuration option.<br><br>Available for all views. |
| Threshold Lines | You can set a threshold for a single metric:<br>• None. You have not set a threshold.<br>• By Symptom Definition. You can set a threshold value based on a symptom definition.<br>• Custom. You can set the threshold value as **Warning**, **Critical**, or **Immediate**. These options are available only for the **Custom** option. |

**Time Settings Tab**

Use the time settings to select the time interval of data transformation.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. Data is calculated based on the browser time.

**Table 219: Time Settings Options**

| Configuration Option | Description |
|---|---|
| Time Range Mode | In Basic mode, you can select date ranges.<br><br>In Advanced mode, you can select any combination of relative or specific start and end dates. |
| Relative Date Range | Select a relative date range of data transformation.<br><br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br><br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br><br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br><br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br><br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br><br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br><br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br><br>Available in Advanced mode. |
| Specific End Date | Select a specific end date of data transformation.<br><br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

**Filter**

The filter option allows you to add additional criteria when the view displays too much information.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**. You can activate Business Hours for the metric selected.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

**Table 220: Filter Add Options**

| Option | Description |
|---|---|
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric or property, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br><br>For instance metrics or properties, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, Sum, and Timestamp. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

**Previous, Next, Create, and Cancel Options**

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

**Distribution View**

Distribution views provide aggregated data about resource distribution in the monitored environment. When you add a distribution type of View to a dashboard, you can click a section of the pie chart or on one of the bars in the bar chart to view the list of objects filtered by the selected segment.

**Where You Find the Distribution View**

From the left menu, click **Operations › Views**. From the **Views** panel, click **Create**. Click **Distribution** from the right panel.

**Name and Configuration Tab**

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Configuration** | |
| Visualization | You can view the data as a pie chart, a bar chart, or a donut chart. When you add a distribution type of View to a dashboard, you can click a section of the pie chart, or on one of the bars in the bar chart, or a section of the donut chart to view the list of objects filtered by the selected segment. You can select the display colors for single or multi-colored charts. |
| Coloring | |
| Colorize | The colors of the slices in the pie chart are displayed in the order of the colors in the color palette. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Select Color | Select the color that you want the chart to appear in. If there is more than one slice in a pie chart, the colors are chosen sequentially from the color palette. In a bar chart, the bars are all the same color. |
| Distribution Type > Dynamic Distribution | |
| Buckets Count | The number of buckets to use in the data distribution. |
| Buckets Size Interval | The bucket size is determined by the defined interval divided by the specified number of buckets. |
| Buckets > Size > Logarithmic bucketing | The bucket size is calculated to logarithmically increasing sizes. This provides a continuous coverage of the whole range with the specified number of buckets. The base of the logarithmic sizing is determined by the given data. |
| Buckets > Size > Simple Max/Min bucketing | The bucket size is divided equally between the measured min and max values. This provides a continuous coverage of the whole range with the specified number of buckets. |
| Distribution Type > Manual Distribution | Specify the number of buckets and the minimum and maximum values of each bucket.<br>You can also select a color for each defined bucket that you specify. |
| Distribution Type > Discrete Distribution | Specify the number of buckets in which VMware Aria Operations distributes the data.<br>If you increase the number of buckets, you can see more detailed data. |
| Distribution Type > Summary | |
| **Settings** | |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br><br>For example, you have a trend view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

## Data Tab

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria OperationsVMware Cloud Foundation Operations collects, calculates, and presents the information for the view.

## How to Add Data to a View

If you selected more than one subject, click the subject for which you add data. Double-click either a metric or a property from the tree in the left panel to add it to the view. For each subject that you select, the data available to add might be different. The Data, Transformation, and Configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
|---|---|
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |
| Data selection tree | Double-click to select a metric or a property. |
| Data column | Click the metric or property to enter configuration details in the Configuration column. |
| Transformation column | Displays the type of transformation applied to the data. |
| Configuration column | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand (MHz) from the **Units** drop-down menu, you can change the value to Hz, KHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |
| Sort order | Orders the values in ascending or descending order. |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>• Minimum. The minimum value of the metric over the selected time range.<br>• Maximum. The maximum value of the metric over the selected time range.<br>• Average. The mean of all the metric values over the selected time range.<br>• Sum. The sum of the metric values over the selected time range.<br>• First. The first metric value for the selected time range.<br>• Last. The last value of a metric within the selected time range.<br>If you have selected **Last** as the transformation in versions before vRealize Operations 6.7, and the end of specified time range is not before the last five minutes, use the **Current** transformation.<br>• Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>• Standard Deviation. The standard deviation of the metric values.<br>• Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>• Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>• Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, `sum/((max + min)/2)`. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select `Sum` as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.<br><br>This option is applicable to the Transformation configuration option. |

**Time Settings Tab**

Use the time settings to select the time interval of data transformation.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. Data is calculated based on the browser time

**Table 221: Time Settings Options**

| Configuration Option | Description |
|---|---|
| Time Range Mode | In Basic mode, you can select date ranges.<br><br>In Advanced mode, you can select any combination of relative or specific start and end dates. |
| Relative Date Range | Select a relative date range of data transformation.<br><br>Available in Basic mode. |

*Table continued on next page*

*Continued from previous page*

| Configuration Option | Description |
|---|---|
| Specific Date Range | Select a specific date range of data transformation.<br><br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br><br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br><br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br><br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br><br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br><br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br><br>Available in Advanced mode. |
| Specific End Date | Select a specific end date of data transformation.<br><br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

**Filter Tab**

The filter option allows you to add additional criteria when the view displays too much information.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**. You can activate Business Hours for the metric selected.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

**Table 222: Filter Add Options**

| Option | Description |
|---|---|
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric or property, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br><br>For instance metrics or properties, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, Sum, and Timestamp. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

**Previous, Next, Create, and Cancel Options**

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

**Text View**

Text views allows you to insert provided text. The text can be dynamic and contain metrics and properties. You can format text to increase or decrease the font size, change the font color, highlight text, and align text to the left, right, or center. You can also make the selected text appear bold, in italics, or underlined. By default the text view is available only for report template creation and modification. You can change this in the **Visibility** option in the **Name and Configuration** tab.

**Where You Find the Text View**

From the left menu, click **Operations › Views**. From the **Views** panel, click **Create**. Click **Text** from the right panel.

**Name and Configuration Tab**

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. This is the default option. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br><br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

## Data Tab

The data definition process includes adding properties, metrics, policies, or data that adapters provide to a view. These are the items by which VMware Aria OperationsVMware Cloud Foundation Operations collects, calculates, and presents the information for the view.

## How to Add Data to a View

If you selected more than one subject, click the subject for which you add data. Double-click the data from the tree in the left panel to add it to the view. For each subject, the data available to add might be different. The data and configuration details are displayed.

You can see a live preview of the view type when you select a subject and associated data, and then click **Select preview source**.

| Option | Description |
|---|---|
| Add Subject | Select the base object type for which the view shows information. The subject you specify determines where the view is applicable. If you select more than one subject, the view is applicable for each of them. |
| **Data Grid options** | |
| Data selection tree | Select a metric or property. |
| Data column | Click the metric or property to enter configuration details in the Configuration column.<br>You can also enter text to display in the view. |
| **Configuration column** | |
| Metric name | Default metric name. |
| Metric label | Customizable label as it appears in the view or report. |
| Units | Depends on the added metric or property. You can select in what unit to display the values. For example, for CPU\| Demand(MHz) from the **Units** drop-down menu, you can change the value to Hz, KHz, or GHz. If you select **Auto**, the scaling is set to a meaningful unit. |
| Transformation | Determines what calculation method is applied on the raw data. You can select the type of transformation:<br>• Minimum. The minimum value of the metric over the selected time range.<br>• Maximum. The maximum value of the metric over the selected time range.<br>• Average. The mean of all the metric values over the selected time range.<br>• Sum. The sum of the metric values over the selected time range.<br>• First. The first metric value for the selected time range.<br>• Last. The last value of a metric within the selected time range.<br>If you have selected **Last** as the transformation in versions before vRealize Operations 6.7, and the end of specified time range is not before the last five minutes, use the **Current** transformation. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise it is null.<br>• Standard Deviation. The standard deviation of the metric values.<br>• Metric Correlation. Displays the value when another metric is at the minimum or maximum. For example, displays the value for memory.usage when cpu.usage is at a maximum.<br>• Forecast. Performs a regressive analysis and predicts future values. Displays the last metric value of the selected range.<br>• Percentile. Calculates the specified percentile for the data range. For example, you can view the 95th percentile, 99th percentile, and so on.<br>• Expression. Allows you to construct a mathematical expression over existing transformations using minus, plus, multiplication, division, unary minus, unary plus, and round brackets. For example, `sum/((max + min)/2)`. You can use the operands of some of the existing transformations such as, `max, min, avg, sum, first, last, current`. You cannot use `standard deviation, forecast, metric correlation, and percentile`.<br>You can customize the metric unit label when you select the Expression transformation. For example, some of the metric units available are, vCPUs, Bps, KBps, Mbps, and MBps. |
| Series Roll up | The time interval at which the data is rolled up. You can select one of the available options. For example, if you select Sum as a Transformation and 5 minutes as the roll-up interval, then the system selects 5-minute interval values and adds them.<br><br>This option is applicable to the Transformation configuration option. |

**Time Settings**

Use the time settings to select the time interval of data transformation. These options are available for all view types, except Image.

You can set a time range for a past period or set a future date for the end of the time period. When you select a future end date and no data is available, the view is populated by forecast data. Data is calculated based on the browser time.

**Table 223: Time Settings Options**

| Configuration Option | Description |
|---|---|
| Time Range Mode | In Basic mode, you can select date ranges.<br><br>In Advanced mode, you can select any combination of relative or specific start and end dates. |
| Relative Date Range | Select a relative date range of data transformation.<br><br>Available in Basic mode. |
| Specific Date Range | Select a specific date range of data transformation.<br><br>Available in Basic mode. |
| Absolute Date Range | Select a date or time range to view data for a time unit such as a complete month or a week. For example, you can run a report on the third of every month for the previous month. Data from the first to the end of the previous month is displayed as against data from the third of the previous month to the third of the current month.<br><br>The units of time available are: **Hours**, **Days**, **Weeks**, **Months**, and **Years**.<br><br>The locale settings of the system determine the start and end of the unit. For example, weeks in most of the European countries begin on Monday while in the United States they begin on Sunday.<br><br>Available in Basic mode. |
| Relative Start Date | Select a relative start date of data transformation.<br><br>Available in Advanced mode. |
| Relative End Date | Select a relative end date of data transformation.<br><br>Available in Advanced mode. |
| Specific Start Date | Select a specific start date of data transformation.<br><br>Available in Advanced mode. |
| Specific End Date | Select a specific end date of data transformation.<br><br>Available in Advanced mode. |
| Currently selected date range | Displays the date or time range you selected. For example, if you select a specific date range from 5/01/2016 to 5/18/2016, the following information is displayed: `May 1, 2016 12:00:00 AM to May 18, 2016 11:55:00 PM`. |

**Filter**

The filter option allows you to add additional criteria when the view displays too much information.

To add a filter to a view, from an existing or new view dialog box, click the **Filter** tab. Fill in the details for each row and click **Add**.

Each subject has a separate filter box. For Alerts Roll up, Alert, and Symptom subjects not all applicable metrics are supported for filtering.

**Table 224: Filter Add Options**

| Option | Description |
|---|---|
| Add | Adds another criteria to the criteria set. The filter returns results that match all the specified criteria.<br>If you add a filter for an instance metric or property, all the instances of the object for which the criteria is met, will be displayed in the preview screen.<br><br>For instance metrics or properties, you can filter based on transformations such as, Current, Average, First, Last, Maximum, Minimum, Sum, and Timestamp. |
| Add another criteria set | Adds another criteria set. The filter returns results that match one criteria set or another. |

**Previous, Next, Create, and Cancel Options**

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

**Image View**

Image views allow you to insert a static image. By default the image view is available only for report template creation and modification.

**Where You Find the Image View**

From the left menu, click **Operations › Views**. From the **Views** panel, click **Create**. Click **Image** from the right panel.

**Name and Configuration Tab**

| Option | Description |
|---|---|
| Name | Name of the view as it appears on the Views page. |
| Description | Description of the view. |
| **Settings** | |
| Make the view available for > Dashboards through the View widget | Select if you want to make the view available in a dashboard. |
| Make the view available for > Report template creation and modification | Select if you want to make the view available in a report template. This is the default option. |
| Make the view available for > Details tab in the environment | Select if you want to make the view available in the **Detail** tab of a specific object. |
| Hide the view for the selected object types | Select an object type for which you do not want to see this view.<br><br>For example, you have a list view with the subject <virtual machines>. It is visible when you select any of its parent objects. You add a data center from the list. The view is not visible anymore on the data center level. |

#### Data Tab

#### How to Add Data to a View

From the **Data** tab, browse and select an image from the left panel to add it to the view.

You can see a live preview of the view type when you select an image and then click **Preview source**.

#### Previous, Next, Create, and Cancel Options

At the end of each tab, you can go to the previous or next tab. You can also cancel the creation of the view. After you have added all the details, click **Create** to create the view.

## Editing, Cloning, and Deleting a View

You can edit, clone, and delete a view. Before you do, familiarize yourself with the consequences of these actions.

#### Edit a View

When you edit a view, all changes are applied to the report templates that contain it. To edit a view, from the left menu, click **Operations › Views**. From the **Views** panel, click **Manage**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Edit**.

#### Clone a View

When you clone a view, the changes that you make to the clone do not affect the source view. To clone a view, from the left menu, click **Operations › Views**. From the **Views** panel, click **Manage**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Clone**.

#### Delete a View

When you delete a view, it is removed from all the report templates that contain it. To delete a view, from the left menu, click **Operations › Views**. From the **Views** panel, click **Manage**. Select a view from the **Views** page, click the vertical ellipsis against the view and select **Delete**.

## Share a View

You might want to present insights into your application or infrastructure performance to a wider audience, using tools such as Microsoft Power BI. By using the view share functionality, you can export a view to an external tool to ensure relevant stakeholders have the latest insights. You can use the provided API end point URL to import view data in a reporting tool. You can also export the configuration details of the view.

#### Where You Can Access the Option to Share Views

From the left menu, click **Operations › Views**. From the **Views** panel, click **Manage**. From the **Views** page on the right, click an existing view, select a preview source and then click the **Share View** icon in the top-right corner.

**Table 225: Options in the Share View Dialog Box**

| Option | Description |
|---|---|
| URL Tab | You can export view data using a URL. The resulting view data includes view configuration data, objects, and metrics metadata, per object metric data, summary and calculated data, and other view specific data. View data usually corresponds to the data shown in the View UI for the specified view configuration. You can copy and paste |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | the URL in other reporting tools for integration with VMware Aria OperationsVMware Cloud Foundation Operations.<br>• Preview Source: You can select the last previewed source.<br>• URL: You can copy the URL for the selected configured view. To use the URL as an authorized user, you need an authorization token.<br>**NOTE**<br>The URL provided is the required REST call for fetching data. See the API Programming Guide for details.<br>• Authorization: The authorization token is generated. To generate the authorization token, enter the **Username** and **Password** and click **Generate**. |
| Export Tab | You can download and export the view configuration details. |

**Table 226: Options in the Share View Dialog Box**

| Option | Description |
|---|---|
| URL Tab | You can export view data using a URL. The resulting view data includes view configuration data, objects, and metrics metadata, per object metric data, summary and calculated data, and other view specific data. View data usually corresponds to the data shown in the View UI for the specified view configuration. You can copy and paste the URL in other reporting tools for integration with VMware Aria OperationsVMware Cloud Foundation Operations.<br>• Preview Source: You can select the last previewed source.<br>• URL: You can copy the URL for the selected configured view. To use the URL as an authorized user, a CSP authentication token is required, you must generate a token separately on the CSP page. To generate a CSP authentication token, follow these steps:<br>1. Log in to the VMware Cloud Services, select your user profile in the top-right corner, and click **My Account**.<br>2. In the **My Account** page, click **API Tokens**, and then click **Generate a New API Token**.<br>3. Select the required organization roles and the service roles. Depending on your requirement, you can specifically select either the organization roles or the service roles.<br>4. Click **Generate**. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | 5. Copy or save the generated token.<br><br>**NOTE**<br>The URL provided is the required REST call for fetching data. See the API Programming guide for details. |
| Export Tab | You can download and export the view configuration details. |

## Including Deleted VMs in List View

In VMware Aria OperationsVMware Cloud Foundation Operations, you can view the deleted objects and the relationship of the objects in the list view. The objects can be VMs, deployments, projects, vApps, and edge gateways. You can also retain the relationship of the objects even after the objects are deleted from the system. The cost of the deleted virtual machines (VMs) is available until the retention period for that VM is over.

### Where You Find Global Settings for Deleted VMs

To specify for how long you want to retain the deleted virtual machines in VMware Aria OperationsVMware Cloud Foundation Operations, from the left menu, click **Administration › Global Settings**. Navigate to **Data Retention** > **Deleted Objects**.

You can also specify the **Deletion Scheduling Interval** which specifies the number of hours between resource deletion scheduling.

Select **Object Deletion Schedule › Add Object Deletion Schedule** and select the virtual machine object from the **Object Kind** drop-down menu, specify the value, and click **Save**. The global setting value for the deleted virtual machine is updated in VMware Aria OperationsVMware Cloud Foundation Operations.

For VMware Aria Automation, the price of the deleted VMs or deployments is added to the corresponding project object as a separate metric. If the deleted VM from VMware Aria Automation is associated with a cost-based pricing policy, then the price for that VM is not added to the corresponding project.

For vCloud Director, the price of deleted VMs, vApps, and Edge Gateways is added to the corresponding organization VDC object again as a separate metric. For vCenter Server, if VM is on unclustered Host, then deleted VM price is assigned to the Host, otherwise to the Cluster.

### How to Include Deleted VMs in List View

The deleted VMs are visible from **Operations › Configurations**, and then click the **Inventory Management** tile. Navigate to **Collection States › Not Existing**.

## User Scenario: Create, Run, Export, and Import a VMware Aria OperationsVMware Cloud Foundation Operations View for Tracking Virtual Machines

As a virtual infrastructure administrator, you use VMware Aria OperationsVMware Cloud Foundation Operations to monitor several environments. You must know the number of virtual machines on each vCenter instance. You define a view to gather the information in a specific order and use it on all VMware Aria OperationsVMware Cloud Foundation Operations environments.

Verify that you have the necessary access rights to perform this task. Your VMware Aria OperationsVMware Cloud Foundation Operations administrator can tell you which actions you can perform.

You will create a distribution view and run it on the main VMware Aria OperationsVMware Cloud Foundation Operations environment. You will export the view and import it in another VMware Aria OperationsVMware Cloud Foundation Operations instance.

**Create a VMware Aria OperationsVMware Cloud Foundation Operations View to view VM Memory Overhead**

To collect and display VM memory overhead vCenter, you create a custom view.

1. From the left menu, click **Operations › Views**.
2. From the **Views** panel, click **Create** and then click **Distribution** from the right panel.
3. From the **New View** dialog box, in the **Name & Configuration** tab, enter `Virtual Machines Memory Overhead`, as the name for the view.
4. Enter a meaningful description for the view.
   For example, `A view showing the VM memory overhead`.
5. From the **Configuration › Visualization** drop-down menu, select **Pie Chart**.
6. From the **Configuration › Distribution Type** option, select **Discrete distribution**.

   Leave **Max number of buckets** deselected because you do not know the number of hosts on each vCenter instance. If you specify a number of buckets and the hosts are more than that number, one of the slices shows unspecified information labeled Others.
7. Click **Next**.
8. From the **Data** tab, click **Subject** to select the object type that applies to the view.
   a) From the drop-down menu, select **Host System**.
9. From the tree panel in the left, in the filter text box enter `VM Overhead`.
10. Double-click **Memory** > `VM Overhead` to add the metric.
11. Retain the default metric configurations and click **Create**.

**Run a View**

To verify the view and capture a snapshot of information at any point, you run the view for a specific object.

Verify that you have the necessary access rights to perform this task. Your VMware Aria OperationsVMware Cloud Foundation Operations administrator can tell you which actions you can perform.

1. From the left menu, click **Inventory**.
2. From the **Inventory** detailed view panel, navigate to a vCenter instance and click the **Details › Views** tab.
   All listed views are applicable for the vCenter instance.
3. From the **Filters** drop-down menu, select **Type › Distribution**.
   You filter the views list to show only distribution type views.
4. Navigate to and click the **Virtual Machines Distribution** view.
   The bottom pane shows the distribution view with information about this vCenter. Each slice represents a host and the numbers on the far left show the number of virtual machines.

**Export a View**

To use a view in another VMware Aria OperationsVMware Cloud Foundation Operations instance, you export a content definition XML file.

Verify that you have the necessary access rights to perform this task. Your VMware Aria OperationsVMware Cloud Foundation Operations administrator can tell you which actions you can perform.

If the exported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

1. From the left menu, click **Operations › Views**.
2. From the **Views** panel, click **Manage**.
3. Select a view and click **Export** from the vertical ellipsis next to the selected view.

### Import a View

To use views from other VMware Aria OperationsVMware Cloud Foundation Operations environments, you import a content definition XML file.

Verify that you have the necessary access rights to perform this task. Your VMware Aria OperationsVMware Cloud Foundation Operations administrator can tell you which actions you can perform.

1. From the left menu, click **Operations › Views**.
2. From the **Views** panel, click **Manage**.
3. Select a view and click the **Import** option from the horizontal ellipsis next to the **Add** option.
4. Browse to select the Virtual Machines Distribution content definition XML file and click **Import**.

   If the imported view contains custom created metrics, such as what-if, supermetrics, or custom adapter metrics, you must recreate them in the new environment.

   **NOTE**
   The imported view overwrites if a view with the same name exists. All report templates that use the existing view are updated with the imported view.

## Configuring Super Metrics

The super metric is a mathematical formula that contains one or more metrics or properties. It is a custom metric that you design to help track combinations of metrics or properties, either from a single object or from multiple objects. If a single metric does not inform you about the behavior of your environment, you can define a super metric.

After you define it, you assign the super metric to one or more object types. This action calculates the super metric for the objects in that object type and simplifies the metrics display. For example, you define a super metric that calculates the average CPU usage on all virtual machines, and you assign it to a cluster. The average CPU usage on all virtual machines in that cluster is reported as a super metric for the cluster.

When the super metric attribute is activated in a policy, you can also collect super metrics from a group of objects associated with a policy.

Because super metric formulas can be complex, plan your super metric before you build it. The key to creating a super metric that alerts you to the expected behavior of your objects is knowing your own enterprise and data. Use this checklist to help identify the most important aspects of your environment before you begin to configure a super metric.

**Table 227: Designing a Super Metric Checklist**

| | |
|---|---|
| Determine the objects that are involved in the behavior to track. | When you define the metrics to use, you can select either specific objects or object types. For example, you can select the specific objects VM001 and VM002, or you can select the object type virtual machine. |
| Determine the metrics to include in the super metric. | If you are tracking the transfer of packets along a network, use metrics that refer to packets in and packets out. In another common use of super metrics, the metrics might |

*Table continued on next page*

*Continued from previous page*

| | |
|---|---|
| | be the average CPU usage or average memory usage of the object type you select. |
| Decide how to combine or compare the metrics. | For example, to find the ratio of packets in to packets out, you must divide the two metrics. If you are tracking CPU usage for an object type, you might want to determine the average use. You might also want to determine what the highest or lowest use is for any object of that type. In more complex scenarios, you might need a formula that uses constants or trigonometric functions. |
| Decide where to assign the super metric. | You define the objects to track in the super metric, then assign the super metric to the object type that contains the objects being tracked. To monitor all the objects in a group, activate the super metric in the policy, and apply the policy to the object group. |
| Determine the policy to which you add the super metric. | After you create the super metric, you add it to a policy. For more information, refer to Policy Workspace in . |

## What Else Can You Do with Super Metrics

- To see the super metrics in your environment, generate a system audit report. For more information, refer to System Audit for .
- To see the super metrics in your environment, generate a system audit report. For more information, refer to the System Audit section in the Information Center.
- To create alert definitions to notify you of the performance of objects in your environment, define symptoms based on super metrics. For more information, refer to Symptom Definitions in .
- Learn about the use of super metrics in policies. For more information, refer to Policy Workspace in .
- Use OPS CLI commands to import, export, configure, and delete super metrics. For more information, refer to the OPS CLI documentation.
- To display metric-related widgets, create a custom set of metrics. You can configure one or more files that define different sets of metrics for a particular adapter and object types. This ensures that the supported widgets are populated based on the configured metrics and selected object type.

## Create a Super Metric

Create a super metric when you want to check the health of your environment, but you cannot find a suitable metric to perform the analysis.

1. From the left menu, click **Operations › Configuration**, and then click the **Super Metrics** tile.
2. Click **Add** .
   The **Create Super Metric** wizard opens.
3. To edit a super metric, click the vertical ellipsis next to the super metric and select **Edit**. You can also edit the super metric using the **EDIT** option in super metrics page.
4. Enter a meaningful name for the super metric such as `Worst VM CPU Usage (%)` in the **Name** text box.

   **NOTE**
   It is important that you have an intuitive name as it appears in dashboards, alerts, and reports. For meaningful names, always use space between words so that it is easier to read. Use title case for consistency with the out of the box metrics and add the unit at the end.

5. Provide a brief summary of the super metric in the **Description** text box.

**NOTE**
Information regarding the super metric, like why it was created and by whom can provide clarity and help you track your super metrics with ease.

6. From the **Object Types** drop-down list, select the object to associate with the super metric and click **Next**.
7. Create the formula for the super metric.

   For example, to add a super metric that captures the average CPU usage across all virtual machines in a cluster, perform the following steps.

   a) Select the function or operator. This selection helps combine the metric expression with operators and/or functions. In the super metric editor, enter `avg` and select the **avg** function.

      You can manually enter functions, operators, objects, object types, metrics, metrics types, property, and properties types in the text box and use the suggestive text to complete your super metric formula. Alternatively, select the function or operator from the **Functions** drop-down menu.

   b) To create a metric expression, enter `Virtual` and select **Virtual Machine** from the object type list.
   c) Add the metric type, enter `usage`, and select the **CPU|Usage (%)** metric from the metric type list.

      **NOTE**
      The expression ends with depth=1 by default. If the expression ends with depth=1, that means that the metric is assigned to an object that is one level above virtual machines in the relationship chain. However, since this super metric is for a cluster which is two levels above virtual machine in the relationship chain, change the depth to 2.
      The depth can also be negative, this happens when you need to aggregate the parents of a child object. For example, when aggregating all the VMs in a datastore, the metric expression ends with depth=-1, because VM is a parent object of datastore. But, if you want to aggregate all the VMs at a Datastore Cluster level, you need to implement 2 super metrics. You cannot directly aggregate from VM to Datastore Cluster, because both are parents of a datastore. For a super metric to be valid, depth cannot be 0 (-1+1=0). Hence, you need to create the first super metric (with depth=-1) for the aggregate at the datastore level, and then build the second super metric based on the first (with depth = 1).

      The metric expression is created.

   d) To calculate the average CPU usage of powered on virtual machines in a cluster, you can add the `where` clause. Enter `where=""`.

      **NOTE**
      The `where` clause cannot point to another object, but can point to a different metric in the same object. For example, you cannot count the number of VMs in a cluster with the CPU contention metric > SLA of that cluster. The phrase "SLA of that cluster " belongs to the cluster object, and not to the VM object. The right operand must also be a number and cannot be another super metric or variable.

   e) Position the pointer between the quotation marks, enter **Virtual**, and select the **Virtual Machine** object type and the **System|Powered ON** metric type.
   f) To add the numeric value for the metric, enter `==1`.
   g) To view hints and suggestions, click **ctrl+space** and select the objects, object types, metrics, metrics types, property, and properties types to build your super metric formula.
   h) Select **This** option from the drop-down menu.

      If **This** option is selected during the creation of a metric expression, it means that the metric expression is associated to the object for which the super metric is created.

8. Select the unit of the super metrics from the **Unit** drop-down.

   **NOTE**
   The super metrics unit configured here can be changed in the metrics charts, widgets, and views.

9. Click **Validate** to verify that the super metric formula has been created correctly.

   After you click the **Preview** button the system selects a random object and displays a metric graph showing values for the current super metric. For example, if you have selected `Host System` in the **Object Types** tab, after you click the **Preview** button it will randomly select a host system object from the list of the available objects and displays the graph for the selected host. Alternatively, you can also type in the object name in the **Object** text box, and the result will also depend on the pre-selected object type.

   a) Expand the **Preview** section.
   A metric graph is displayed showing values of the metric collected for the object. Verify that the graph shows values over time.

   b) Click **Next**.

   The Policies page is displayed.

10. Select the policy which you want to associate with super metric and click **Update**.
    The selected policy is applied to the super metric. You can view the super metric you created, the associated object type, and policy on the **Super Metrics** page.

## Enhancing Your Super Metrics

You can enhance your super metrics by using clauses and resource entry aliasing.

### Where Clause

The `where` clause verifies whether a particular metric value can be used in the super metric. Use this clause to point to a different metric of the same object, such as `where=(${metric=metric_group|my_metric} > 0)`.

For example: `count(${objecttype = ExampleAdapter, adaptertype = ExampleObject, metric = ExampleGroup|Rating, depth=2, where =($value==1})`

### IsFresh Function

Use the `isFresh` function in the `where` clause to check if the last value of the metrics is fresh or not.

For every metric published in VMware Aria OperationsVMware Cloud Foundation Operations, the point with the latest publishing time is called as the last point of that metric. The value of that metric's last point is called the last value of that metric. A metric's last point is considered fresh when the time elapsed after the metric's last point is lesser than the estimated publishing interval of that metric.

The `isFresh` function returns true if the last value of the metrics is fresh. For example, in the following scenarios, the function:

- `${this, metric=a|b, where=($value.isFresh())}`, returns the last value of the metric a|b if the last value is fresh.
- `${this, metric=a|b, where=($value == 7 && $value.isFresh())}`, returns the last value of the metric a|b if it is equal to seven and is fresh.
- `${this, metric=a|b, where=(${metric=c|d} == 7 && ${metric=c|d}.isFresh())}`, returns the last value of the metric a|b only if the last value of the metric c|d is equal to seven and is fresh.

**Resource Entry Aliasing**

Resource entries are used to retrieve metric data from VMware Aria OperationsVMware Cloud Foundation Operations for computing super metrics. A resource entry is the part of an expression which begins with `$` followed by a `{..} block`. When computing a super metric, you might have to use the same resource entry multiple times. If you have to change your computation, you must change every resource entry, which might lead to errors. You can use resource entry aliasing to rewrite the expression.

The following example, shows a resource entry that has been used twice.

```
(min(${adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=0)}) + 0.0001)/(max(${adaptertype=VMWARE,
objecttype=HostSystem, attribute=cpu|demand|active_longterm_load, depth=5,
where=($value>=0)}) + 0.0001)"
```

The following example shows how to write the expressing using resource entry aliasing. The output of both expressions is the same.

```
(min(${adaptertype=VMWARE, objecttype=HostSystem, attribute= cpu|demand|
active_longterm_load, depth=5, where=($value>=0)} as cpuload) + 0.0001)/(max(cpuload) +
0.0001)"
```

Follow these guidelines when you use resource entry aliasing:

- When you create an alias, make sure that after the resource entry you write `as` and then `alias:name`. For example: **${…} as alias_name**.
- The alias cannot contain the ()[]+-*/%|&!=<>,.?:$ special characters, and cannot begin with a digit.
- An alias name, like all names in super metric expressions, is case-insensitive.
- Use of an alias name is optional. You can define the alias, and not use it in an expression.
- Each alias name can be used only once. For example: `${resource1,…} as r1 + ${resource2,…} as R1.`
- You can specify multiple aliases for the same resource entry. For example: **${…} as a1 as a2**.

**Conditional Expression `?:` Ternary Operators**

You can use a ternary operator in an expression to run conditional expressions.

For example: `expression_condition ? expression_if_true : expression_if_false`.

The result of the conditional expression is converted to a number. If the value is not 0, then the condition is assumed as true.

For example: `-0.7 ? 10 : 20` equals 10. `2 + 2 / 2 - 3 ? 4 + 5 / 6 : 7 + 8` equals 15 (7 + 8).

Depending on the condition, either `expression_if_true` or `expression_if_false` is run, but not both of them. In this way, you can write expressions such as, `${this, metric=cpu|demandmhz} as a != 0 ? 1/a : -1.` A ternary operator can contain other operators in all its expressions, including other ternary operators.

For example: `!1 ? 2 ? 3 : 4 : 5` equals 5.

## Exporting and Importing a Super Metric

You can export a super metric from one VMware Aria OperationsVMware Cloud Foundation Operations instance and import it to another VMware Aria OperationsVMware Cloud Foundation Operations instance. For example, after developing a super metric in a test environment, you can export it from the test environment and import it use in a production environment.

If the super metric to import contains a reference to an object that does not exist in the target instance, the import fails. VMware Aria OperationsVMware Cloud Foundation Operations returns a brief error message and writes detailed information to the log file.

1. Export a super metric.
   a) From the left menu, click **Operations › Configuration**, and then click the **Super Metrics** tile.
   b) Select the super metric to export, click horizontal ellipsis and then click **Export**.
   VMware Aria OperationsVMware Cloud Foundation Operations creates a super metric file, for example, **`SuperMetric.json`**.
   c) Download the super metric file to your computer.
2. Import a super metric.
   a) From the left menu, select **Configure** and then click **Super Metrics**.
   b) Click the horizontal ellipsis and then click **Import**.
   c) (Optional). If the target instance has a super metric with the same name as the super metric you are importing, you can either overwrite the existing super metric or skip the import, which is the default.

## Super Metrics Tab

A super metric is a mathematical formula that contains a combination of one or more metrics for one or more objects. With super metrics you can assess information more quickly when you are observing fewer metrics.

### Where You Configure Super Metrics

From the left menu, click **Operations › Configuration**, and then click the **Super Metrics** tile.

To view the details for a specific super metric, click the super metric from the list. The super metric details are displayed in the right-side panel. The super metric details include the assigned object types, formula and policies activated for the selected super metric.

**Table 228: Configuration Options for Super Metrics**

| Option | Description |
|---|---|
| Toolbar | Use the toolbar selections to manage super metric options.<br>• Add New Super Metric. Starts the Create Super Metric workspace.<br>• Edit Selected Super Metric. Starts the Create Super Metric workspace.<br>• Clone Selected Super Metric. Duplicates the super metric. Edit the clone or associate it with a different object type.<br>• Delete Selected Super Metric.<br>• Export Selected Super Metric. Exports a super metric to use in another VMware Aria OperationsVMware Cloud Foundation Operations instance. See Exporting and Importing a Super Metric.<br>• Import Super Metric. Imports a super metric to this VMware Aria OperationsVMware Cloud Foundation Operations instance. See Exporting and Importing a Super Metric. |
| Super Metrics list | Configured super metrics listed by name and formula description. |

### Enhancements to the Super Metric Functions

In the earlier implementation of aggregate functions in super metrics, you had to explicitly specify the Adapter Kind and Resource Kind in the formula.

Old Formula

```
count(${adaptertype=VMWARE, objecttype=HostSystem,attribute=badge|health, depth=1})
```

The new implementation of aggregate function provides a way to define a super-metric without explicitly specifying the Resource Kind. You can use "objecttype=*" in the super-metric formula which indicates you to consider all Resource Kinds having the specified attribute.

New Formula

```
count(${adaptertype=VMWARE, objecttype=*,attribute=badge|health, depth=1})
```

> **NOTE**
> The explicit specification of "adaptertype" remains mandatory. However, "*" can be used only to select all Resource Kinds for the given Adapter Kind.

## Manage Super Metric Workspace

You use the Manage Super Metric workspace to create or edit a super metric. The toolbar helps you to build the mathematical formula with the objects and metrics you select.

### Where You Configure Super Metrics

From the left menu, click **Operations › Configuration**, and then click the **Super Metrics** tile.

**Table 229: Super Metrics Workspace Options**

| Option | Description |
|---|---|
| Super Metric | • Name. The name you give to the super metric.<br>• Description. The textual description you give about the super metric. |
| Object Types Pane | Use this page to associate the super metric with an object type. You can use this list to select the object type with the metrics to measure. The object type selection affects the list of objects, metrics, and attribute types displayed. |
| Formula | Define the formula you want to associate with the super metric. You can preview and validate the formula before you create it. Use drop-down menu to select the metrics to add to the formula.<br>• Functions. Mathematical functions that operate on a single object or group of objects. See Super Metric Functions and Operators.<br>• Operators. Mathematical symbols to enclose or insert between functions. See Enhancing Your Super Metrics.<br>• This Object. Assigns the super metric to the object selected in the Object pane and displays `this` in the formula instead of a long description for the object. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Unformatted. Slide the toggle button to the right to view the unformatted version of the formula, slide the button to the left to view the formula in textual format.<br>• Preview. Shows the super metric in a graph. Look at the graph so that you can verify that VMware Aria OperationsVMware Cloud Foundation Operations is calculating the super metric for the target objects that you selected. |
| Policies | Displays the policies related to the object types you assigned your super metric to. |

### Super Metric Functions and Operators

VMware Aria OperationsVMware Cloud Foundation Operations includes functions and operators that you can use in super metric formulas. The functions are either looping functions or single functions.

#### Looping Functions

Looping functions work on more than one value.

**Table 230: Looping Functions**

| Function | Description |
|---|---|
| avg | Average of the collected values. |
| combine | Combines all the values of the metrics of the included objects in a single metric timeline. |
| count | Number of values collected. |
| max | Maximum value of the collected values. |
| min | Minimum value of the collected values. |
| sum | Total of the collected values. |

**NOTE**
VMware Cloud Foundation Operations 5.x included two sum functions: **`sum (expr)`** and **`sumN (expr, depth)`**. VMware Cloud Foundation Operations 6.x includes one sum function: **`sum (expr)`**. Depth is set at depth=1 by default. For more information about setting depth, refer to Create a Super Metric.

#### Looping Function Arguments

The looping function returns an attribute or metric value for an object or object type. An attribute is metadata that describes the metric for the adapter to collect from the object. A metric is an instance of an attribute. The argument syntax defines the desired result.

For example, CPU usage is an attribute of a virtual machine object. If a virtual machine has multiple CPUs, the CPU usage for each CPU is a metric instance. If a virtual machine has one CPU, then the function for the attribute or the metric return the same result.

**Table 231: Looping Function Formats**

| Argument syntax example | Description |
|---|---|
| *funct*(${this, metric =*a*\|*b:optional_instance*\|*c*}) | Returns a single data point of a particular metric for the object to which the super metric is assigned. This super metric does not take values from the children or parents of the object. |
| *funct*(${this, attribute=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points for attributes of the object to which the super metric is assigned. This super metric does not take values from the child or parent of the object. |
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, resourcename=*resname*, identifiers={id1=*val1*id2=*val2*,...}, metric=*a*\|*b:instance*\|*c*}) | Returns a single data point of a particular metric for the *resname* specified in the argument. This super metric does not take values from the children or parents of the object. |
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, resourcename=*resname*, identifiers={*id1=val1, id2=val2,*...}, attribute=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points. This function iterates attributes of the *resname* specified in the argument. This super metric does not take values from the child or parent of the object. |
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, depth=*dep*}, metric=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points. This function iterates metrics of the *reskind* specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects, where *depth* describes the object location in the relationship chain.<br>For example, a typical relationship chain includes a data center, cluster, host, and virtual machines. The data center is at the top and the virtual machines at the bottom. If the super metric is assigned to the cluster and the function definition includes depth = 2, the super metric takes values from the virtual machines. If the function definition includes depth = -1, the super metric takes values from the data center. |
| *funct*(${adaptertype=*adaptkind*, objecttype=*reskind*, depth=*dep*}, attribute=*a*\|*b:optional_instance*\|*c*}) | Returns a set of data points. This function iterates attributes of the *reskind* specified in the argument. This super metric takes values from the child (depth > 0) or parent (depth < 0) objects. |

For example, `avg(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|usage_average, depth=1})` averages the value of all metric instances with the cpu|usage_average attribute for all objects of type **`VirtualMachine`** that the vCenter adapter finds. VMware Aria OperationsVMware Cloud Foundation Operations searches for objects one level below the object type where you assign the super metric.

## Single Functions

Single functions work on only a single value or a single pair of values.

**Table 232: Single Functions**

| Function | Format | Description |
|---|---|---|
| *abs* | abs(x) | Absolute value of x. x can be any floating point number. |
| *acos* | acos(x) | Arccosine of x. |
| *asin* | asin(x) | Arcsine of x. |
| *atan* | atan(x) | Arctangent of x. |
| *ceil* | ceil(x) | The smallest integer that is greater than or equal to x. |
| *cos* | cos(x) | Cosine of x. |

*Table continued on next page*

*Continued from previous page*

| Function | Format | Description |
|---|---|---|
| *cosh* | cosh(x) | Hyperbolic cosine of x. |
| *exp* | exp(x) | e raised to the power of x. |
| *floor* | floor(x) | The largest integer that is less than or equal to x. |
| *log* | log(x) | Natural logarithm (base *x*) of x. |
| *log10* | log10(x) | Common logarithm (base 10) of x. |
| *pow* | pow(x,y) | Raises x to the y power. |
| *rand* | rand() | Generates a pseudo random floating number greater than or equal to 0.0 and less than 1.0. |
| *sin* | sin(x) | Sine of x. |
| *sinh* | sinh(x) | Hyperbolic sine of x. |
| *sqrt* | sqrt(x) | Square root of x. |
| *tan* | tan(x) | Tangent of x. |
| *tanh* | tanh(x) | Hyperbolic tangent of x. |

## Operators

Operators are mathematical symbols and text to enclose or insert between functions.

**Table 233: Numeric Operators**

| Operators | Description |
|---|---|
| + | Plus |
| - | Subtract |
| * | Multiply |
| / | Divide |
| % | Modulo |
| == | Equal |
| != | Not equal |
| < | Less than |
| <= | Less than, or equal |
| > | Greater than |
| >= | Greater than, or equal |
| \|\| | Or |
| && | And |
| ! | Not |
| ? : | Ternary operator. If/then/else<br>For example: `conditional_expression ? expression_if_condition_is_true : expression_if_condition_is_false`<br><br>For more information about ternary operators, see Enhancing Your Super Metrics. |
| ( ) | Parentheses |

*Table continued on next page*

*Continued from previous page*

| Operators | Description |
|---|---|
| [ ] | Use in an array of expressions |
| [x, y, z] | An array containing x, y, z. For example, min([x, y, z]) |

**Table 234: String Operators**

| String Operators | Description |
|---|---|
| equals | Returns true if metric/property string value is equal to specified string. |
| contains | Returns true if metric/property string value contains specified string. |
| startsWith | Returns true if metric/property string value starts with the specified prefix. |
| endsWith | Returns true if metric/property string value ends with the specified suffix. |
| !equals | Returns true if metric/property string value is not equal to specified string. |
| !contains | Returns true if metric/property string value does not contain specified string. |
| !startsWith | Returns true if metric/property string value does not start with the specified prefix. |
| !endsWith | Returns true if metric/property string value does not end with the specified suffix. |

**NOTE**
String operators are valid in 'where' condition only. For example: `${this, metric=summary|runtime|isIdle, where = "System Properties|resource_kind_type !contains GENERAL"}`

## About Licenses

VMware Aria OperationsVMware Cloud Foundation Operations provides a centralized plane for license management and consumption visibility. Starting with VMware Aria OperationsVMware Cloud Foundation Operations 8.18, you can manage and monitor the consumption of all available licenses from the registered vCenter systems in VMware Aria OperationsVMware Cloud Foundation Operations.

You can add new VMware Cloud Foundation Solution License keys, VMware vSphere Foundation Solution License keys, or vSAN per TiB license keys and add them to registered vCenter systems in VMware Aria Operations. For more information, see Managing Licenses. You can also view the consumption data and usage trends of the licenses available in registered vCenter systems. For more information, see License Usage Analytics.

vSphere 8 Enterprise Plus for vSphere Foundation entitles VMware Aria OperationsVMware Cloud Foundation Operations to the advanced edition, while vSphere 8 Enterprise Plus for VCF entitles VMware Aria OperationsVMware Cloud Foundation Operations to the enterprise edition.

All the older VMware Aria OperationsVMware Cloud Foundation Operations license types and license-related functionalities are also valid and supported. For more information, see VMware Aria OperationsVMware Cloud Foundation Operations License Keys.

**NOTE**
Starting with VMware Aria Operations 8.16, license consumption for VMware Aria Operations is not required for entitled setups, the license consumption is tracked against your vCenter systems.

A newly deployed VMware Aria Operations works in evaluation mode for 60 days if no valid entitlement is available.

VMware Aria Operations always gives precedence to vCenter cloud accounts with vSphere 8 Enterprise Plus for vSphere Foundation or vSphere 8 Enterprise Plus for VCF licenses, unless a higher edition VMware Aria Operations license is available.

The entitlement is checked every time a new vCenter cloud account is added. For details on how to configure the vCenter cloud account, see Configure a vCenter Cloud Account in VMware Aria OperationsVMware Cloud Foundation Operations.

If the entitlement expires or if VMware Aria Operations is unable to fetch valid entitlement from the connected vCenter systems, the product enters a 90 day grace period. The grace period does not restrict any of the functionalities, but you must add a valid license or entitlement before the grace period ends. After the grace period ends, if there is a valid VMware Aria Operations license available, the product continues to work with the license. If VMware Aria Operations cannot find any valid license or entitlement at the end of the grace period, it enters the restricted mode. In the restricted mode, you can only access the following pages:

- Subscriptions: License Management and Legacy Licenses
- Administration: Integrations and Cloud Proxies

**NOTE**
Data collection does not stop in the restricted mode.

To exit the restricted mode, perform any one of the following actions:

- Add a vSphere 8 Enterprise Plus for vSphere Foundation or a vSphere 8 Enterprise Plus for VCF license from the **Subscriptions › License Management** page or enter a valid VMware Aria Operations license key from the **Subscriptions › Legacy Licenses** page.

  **NOTE**
  After you add a valid license on the License Management page, you must also add the license to a vCenter system and assign it to the vCenter asset by using the vSphere client. VMware Aria Operations can then fetch the valid entitlement and work with the available edition.

- Add a new entitled vCenter cloud account from **Administration › Integrations** or use the suite API.
- Add a valid license to a registered vCenter system using the vSphere client.

After a valid license or entitlement is available, VMware Aria Operations exits the restricted mode and functions as per the edition available with the license or entitlement.

## Managing Licenses

VMware Aria OperationsVMware Cloud Foundation Operations provides unified plane for license management and consumption visibility for all the licenses used by vCenter systems in VMware Aria Operations.

Before you begin, ensure you either have a VMware Aria Operations license key or a vCenter cloud account with vSphere 8 Enterprise Plus for vSphere Foundation or vSphere 8 Enterprise Plus for VCF licenses. You can log in to Broadcom Support Portal to get your licenses.

To start managing your licenses, from the left menu, click **Subscription › License Management**.

**NOTE**
To view the license management page, you must have the **Subscriptions > Licensing > View Licenses** permission. To manage the license management page, you must have the **Subscriptions > Licensing > View Licenses** and **Subscriptions > Licensing > Manage Licenses** permissions.

On the Licenses tab, you can view the license usage data collected from the registered vCenter systems in VMware Aria Operations. You can view the per product usage overview and usage of individual licenses that are available in the VMware Aria Operations inventory. The data on the licenses tab updates after each collection cycle.

- **Usage Overview**: Displays the product-based overview of the licenses. You can view the aggregated data of all your licenses in a single product.
- **Licenses**: Displays all the individual licenses that are collected by VMware Aria Operations or are added to VMware Aria Operations.

**Usage Overview**

Use the **Usage Overview** section to view the aggregated data for licenses per product. Click the **Chart View** or **List View** icon to view the data displayed in a graphical or a list format.

Each product type can have one or more licenses and the data displayed in the usage overview section is a combination of all the licenses in that product type.

**Licenses**

Use the **Licenses** section to view individual license data. The licenses table displays all licenses collected by VMware Aria Operations from the registered vCenter systems. You can add new licenses, add them to vCenter, or remove them from VMware Aria Operations.

You can add a new license and then add the license to vCenter, for more information, see Add Licenses. After you add a license to the VMware Aria Operations inventory, it appears in the licenses table. To view the usage data of a new license key you must add it to a registered vCenter system. You can also select any of the available licenses and add it to vCenter. To delete licenses from the VMware Aria Operations inventory, first remove it from vCenter.

> **NOTE**
> VMware Aria Operations can use the default or the action credentials to communicate with vCenter for licensing purposes. Per vCenter, make sure the credentials you choose have the **Global.Licenses** privilege assigned and are a member of the LicenseService.Administrators Single Sign-On group. If you choose the default credentials, make sure no action credentials are configured.
> If action credentials are provided, those are used to perform the operation, if not, the default credentials are used.

You can view in-depth data of the individual licenses using the **License Details** pane. The assigned assets display all vCenter systems that are linked to the license key along with the usage and last collection time.

> **NOTE**
> The Evaluation Mode license key appears if one or more vCenter systems have assets that work in evaluation mode. To view all vCenter systems with assets that work in evaluation mode, open the **License Details** pane.

### Add Licenses

You can add VMware Cloud Foundation Solution License keys, VMware vSphere Foundation Solution License keys, or vSAN per TiB license keys to the VMware Aria OperationsVMware Cloud Foundation Operations inventory. After you add a license key, you can then add it to registered vCenter systems.

1. From the left menu, click **Subscription › License Management**.
2. Click **Add License** to add a new license.
3. Enter the license key in the **Add License Keys** text box and click **Next**.

> **NOTE**
> You can only add VMware Cloud Foundation Solution License keys, VMware vSphere Foundation Solution License keys, or vSAN per TiB license keys that are new and active. The validation fails if the license key already exists in the VMware Aria Operations inventory or if it is expired.

4. From the list of vCenter systems, select the **vCenter Name** to which you want to add the license keys, and click **Next**.

   vCenter systems that already contain the license key are not listed.

   **NOTE**
   You can add the license to registered vCenter systems in VMware Aria Operations. To assign the license to vCenter systems, you must use the vSphere client.

5. Review and verify your selections and click **Add**.
   The license key is now added to the VMware Aria OperationsVMware Cloud Foundation Operations inventory and appears in the licenses table on the **License Management** page.

You can also add a license to vCenter systems after adding the license key to the VMware Aria Operations inventory. On the License Management page, select the license and click **Add to vCenter**. From the list of available vCenter systems, select the vCenter to which you want to add the license.

## License Usage Analytics

You can view and export the historical usage data of your licenses per product in VMware Aria OperationsVMware Cloud Foundation Operations. Use the usage analytics tab to view the usage overview trend of the product and view the individual license usage trend for each license belonging to the product. You can view the trend for a period ranging between 1 and 90 days. You can view the aggregated usage, overage, and the capacity of the license for the product as per the dates selected in the time range.

1. From the left menu, click **Subscription › License Management**, and then click the **Usage Analytics** tab.
2. Select a **Product** from the drop-down menu.
3. Select the time range, click the calendar icon to select the **From** and **To** dates, and then click **Submit**.
   The **Usage Trend Overview** shows the graphical representation of the license usage of the product in a chronological order. You can view the daily usage detail of the product's licenses by clicking on the bars of the trend and the data is displayed in a tabular format underneath the trend.

   **NOTE**
   Usage trend for products with unlimited capacity licenses is not displayed.

   The **License Usage Trend** displays the usage trends of all the licenses that contribute to the product. You can hover over the graph of each license to view the date, capacity, and usage related data.
4. Click **Export** to export the historical usage data of your products.
   a) Select the **Product** from the drop-dowm menu.
   b) Select the **Start Date** and **End Date** and then click **Export**.

## VMware Aria OperationsVMware Cloud Foundation Operations License Keys

License Keys
To activate VMware Aria OperationsVMware Cloud Foundation Operations monitoring , you add licenses at installation or later. You track licenses so that you know what VMware Aria OperationsVMware Cloud Foundation Operations can monitor and when your licenses expire. The product works in evaluation mode until a new valid license key is installed. After you log in to the user interface of VMware Aria OperationsVMware Cloud Foundation Operations, if you see that you are using an evaluation license, consider applying for a new license before the end of the evaluation period.

Before you begin, ensure you either have a VMware Aria Operations license key or a vCenter cloud account with vSphere 8 Enterprise Plus for vSphere Foundation or vSphere 8 Enterprise Plus for VCF licenses.

**NOTE**
Starting VMware Aria OperationsVMware Cloud Foundation Operations 8.16, license consumption for VMware Aria OperationsVMware Cloud Foundation Operations is not required for entitled setups, the consumption is tracked against your vCenter. For more information, see About Licenses.

## How License Keys Work

License keys activate the solution or product and are available in varying levels. Higher levels typically allow VMware Aria OperationsVMware Cloud Foundation Operations to monitor more objects.

## Where You Find the License Keys

1. From the left menu, click **Subscriptions › Legacy Licenses**.

## License Key Options

The options include toolbar and data grid options.

Click **Add** or click the **Hortizontal Ellipses** to refresh or remove license keys.

**Table 235: License Key Toolbar Options**

| Option | Description |
|---|---|
| Add | Select a solution or product, and then enter and validate a license key for it. |
| Delete | Remove a license key. |
| Refresh License Usage | Update license usage details. |

Use the data grid options to view item details.

**Table 236: License Key Data Grid Options**

| Option | Description |
|---|---|
| Product or Solution | Name of the product or solution associated with the key. |
| License Type | Level of the license.<br><br>To view the license edition, click the  icon, and then click **About**. The About VMware Aria OperationsVMware Cloud Foundation Operations dialog box opens. You can view the version no and the license edition that is in use. |
| License Capacity | Number of objects that the license allows the product to monitor. |
| License Usage | Number of monitored objects that count against the capacity. If you have an unlimited capacity, this number is zero (0). |
| Status | Indicates whether the license is valid. |
| Expiry | Date and time when the license expires. |
| License Information (below) | Details for the selected license key. |
| Overview | Solution or product, expiration, capacity, type, and use of the selected license key. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Associated License Groups | License groups that this key is a member of, and the number of objects in the groups. |

## VMware Aria OperationsVMware Cloud Foundation Operations License Groups

License Groups
Like other VMware Aria OperationsVMware Cloud Foundation Operations groups, you create a license group of objects as a way of gathering those objects for data collection. In this case, you are associating the objects with a product license.

### How License Groups Work

License groups require that you select one or more keys that you already added for solution or product activation, and add objects as members to a custom group for those licenses. You might, for example, want to add objects into groups that are associated with a particular level of license key, and monitor or manage by level of key in order to control licensing costs.

### Where You Find the License Groups

1. From the left menu, click **Subscriptions › Legacy Licenses**, and then click the **License Groups** tab.
2. Click the **License Groups** tab.

### License Groups

**vCloud Suite**

Host CPU-based licenses applied to an object type "Host system" for a given set of clusters. When you apply a CPU license to a group containing Hosts, the VMs on the Hosts will still show "License is invalid" watermark.

**VM Licenses**

VM based licenses applied to an object type "Virtual Machine" for all other VMs except those on hosts licensed with vCloud Suite. When you apply a VM license key to Virtual Machines, the Hosts on which those VMs run will still show the "License is invalid" watermark.

**NOTE**
In VMware Aria OperationsVMware Cloud Foundation Operations, it is possible to mix Operating System Instance (OSI) and CPU based licenses. By mixing different kind of licenses, you will need to perform extra configurations, like creating separate license groups for each type of license keys (one for CPU and one for OSI (VM)). It is recommended that you use non overlapping exclusive Licensing Groups to have the best advantage when you mix OSI (VM) and CPU licensing.
However, in VMware Aria OperationsVMware Cloud Foundation Operations you cannot mix standard edition of VMware Aria OperationsVMware Cloud Foundation Operations with any other advanced and enterprise licenses.

**Dynamic**

Use dynamic membership criteria, not static "Always include/exclude" lists to avoid manual maintenance of license groups.

**NOTE**
When the license is applied to the respective Object type of each License key, the related objects (parent or children) are also going to have to be included in membership for the License Group. License in invalid" watermark appears in VMware Aria OperationsVMware Cloud Foundation Operations 6.6 and later. For more information, see the following KB article 51556.

**License Group Options**

The license group options include toolbar and data grid options.

Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

**Table 237: License Group Toolbar Options**

| Option | Description |
|---|---|
| Add | Launch a wizard to select licenses and objects, to create a new license group. |
| Edit | Launch a wizard to select licenses and objects, to change a license group. |
| Delete | Remove a license group. |

Use the data grid options to view item details.

**Table 238: License Group Data Grid Options**

| Option | Description |
|---|---|
| License Group | Name of the license group |
| Total Members | Number of objects in the license group |
| Licensable Usage | Number of objects in the group that count against the license in order to monitor them. If you have a license for unlimited object monitoring, this number is zero (0). |
| License Group Information (below) | Details for the selected license group |
| Overview | Name, license serial number, and number of keys associated with the selected license group |
| Members | List of objects associated with the selected license group |

## Configuring Administration Settings

After VMware Aria OperationsVMware Cloud Foundation Operations is installed and configured, you can use administration settings to manage your environment. You find most administration settings under the Administration selection of the VMware Aria OperationsVMware Cloud Foundation Operations interface.

## Modifying Global Settings

The global settings control the system settings for VMware Cloud Foundation OperationsVMware Aria Operations, including data retention and system timeout settings. You can modify one or more of the settings to monitor your environment better. These settings affect all your users.

The global settings do not affect metric interactions, color indicators, or other object management behaviors. These behaviors are configured in your policies.

Settings related to managing objects with VMware Cloud Foundation OperationsVMware Aria Operations are available on the **Inventory** page.

## Global Settings Best Practices

Most of the settings pertain to how long VMware Cloud Foundation OperationsVMware Aria Operations retains collected and process data.

The default values are common retention periods. You might need to adjust the time periods based on your local policies or disk space.

### Access Control: Password Policy

To ensure security in VMware Cloud Foundation OperationsVMware Aria Operations, you must manage user passwords. Determine the criteria used for account lockout, password strength, and the password change policy. When a user session becomes inactive for 30 minutes, the session times out, and the user must log in to VMware Cloud Foundation OperationsVMware Aria Operations again.

## Where You Manage the Password Policy

1. From the left menu, click **Administration › Global Settings**.
2. Click **User Access** and navigate to Password Policy.

**Account Lockout**

Indicates whether the account lockout is in effect, and indicates the number of login attempts allowed before the account is locked. The account lockout policy is activated by default.

**Password Strength**

Indicates whether the policy that requires users to strengthen their password is in effect, and the minimum number of characters required to make a strong password. The password strength policy is activated by default.

**Password Change**

Indicates whether the policy that requires users to change their password is in effect, how often the password expires, and whether users will receive a warning. The account password change policy is activated by default.

**Concurrent UI login sessions**

Indicates whether a user can have concurrent UI login. The concurrent UI login sessions policy is activated by default.

**Allow non-imported vIDM user access**

The policy allows non-imported VMware Identity Manager users to be created automatically as read-only users upon first access. If deactivated, only VMware Identity Manager imported users or users belonging to imported VMware Identity Manager groups will be granted access.

## Modify the Password Policy Settings

You can modify the following password policy settings .

**Table 239: Access Control Edit Password Policy Settings**

| Option | Description |
|---|---|
| Account Lockout | Modify the settings to lock user accounts. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Activate Account Lockout Policy. Activate the policy to lock user accounts. For a super administrator user, the account lockout policy is activated by default and cannot be deactivated. The super administrator user account is locked for approximately one hour, and then unlocked.<br>• Number of failed login attempts before lockout. Indicates the number of tries that a user can attempt to log in to VMware Cloud Foundation OperationsVMware Aria Operations before their account is locked. The default number of tries is seven.<br>• Login fail eviction time. Indicates the time available for the user to complete the login attempts. You can set fail evication time in seconds. |
| Password Strength | Modify the settings required for users to create strong passwords.<br>• Activate Password Strength Policy. When selected, activates the policy to require users to strengthen their password.<br>• Minimum password length. Indicates the number of characters required for user passwords. The default length is eight characters.<br>• Passwords must contain numbers. When selected, users must include a combination of numbers and letters in their passwords.<br>• Passwords must not match user names. To ensure security, users are not allowed to use their user name as their password.<br>• Passwords must contain at least one uppercase and one lowercase letter. When selected, users must include one or more uppercase and lowercase characters.<br>• Passwords must contain special characters. When selected, users must include one or more special characters. Special characters include: !@#$%^&*+= |
| Password Change | Modify the settings required for users to change their password.<br>• Activate Password Change Policy. Activate the policy to require users to change their password at specific intervals.<br>• Passwords expire every 90 days. Users receive notification five days before the password expires.<br>• Warn users 5 days prior to expiration. Indicate when to have VMware Cloud Foundation OperationsVMware Aria Operations notify users that their password will expire. The default is five days before their password expires. |

**Access Control: Login Message**

To provide support for Security Technical Implementation Guide (STIG), you can add a Standard Mandatory DoD Notice and Consent Banner for the users who access VMware Aria OperationsVMware Cloud Foundation Operations. Use the login message global setting to set a message that requires an explicit consent before logging in to VMware Cloud Foundation OperationsVMware Aria Operations.

1. From the left menu, click **Administration › Global Settings**.
2. Click **User Access** and navigate to Login Message.
3. To activate the login message, slide the **Activated** bar to the right. The **Display Content** dialog is displayed.
4. Enter the **Title** and enter the content you want to display.

   **NOTE**
   You can add text and images copied from an external source and edit it using the formatting options available.

5. Enter the button label for users to click to provide their consent. The label **Agree** is entered by default.
6. Click **Preview** to view how the message will appear on the Login screen.

7. Click **Save**.

### Access Global Settings

With global settings, you set times to delete objects, set timeouts, store historical data, use dynamic threshold and capacity calculations, and determine how vCenter users log in. For automated actions, you can select whether to allow actions to be triggered from alert recommendations automatically.

1. From the left menu, click **Administration › Global Settings**.
2. The global settings are categorised into Data Retention, Cost/Price, User Access, System Settings. Select one of the categories to edit the global settings.
3. If you cannot find a settings, type the name of the setting in the search bar.
4. When you edit a setting, **Save** and **Cancel** buttons appear. Make your selection to complete the editing process.

### List of Global Settings

The global settings determine how VMware Cloud Foundation OperationsVMware Aria Operations retains data, keeps connection sessions open, and other settings. These are system settings that affect all users. Some of these settings are not editable. The global settings are grouped into four types. .

#### Settings in the Data Retention Category

**Table 240: Global Setting Default Values and Descriptions**

| Setting | Default Value | Description |
|---|---|---|
| Action History | 14 days<br><br>30 days | Number of days to retain the recent task data for actions.<br><br>The data is purged from the system after the specified number of days.<br><br>You can edit this setting from a minimum of one day to a maximum of 30 days in VMware Cloud Foundation Operations. |
| Deleted Objects | 168 hours<br><br>144 hours | This setting determines the number of hours for which objects in the inventory must be retained when:<br>• the objects are deleted from an adapter data source or server before deleting them from VMware Cloud Foundation OperationsVMware Aria Operations<br>• A cloud account is deleted from VMware Cloud Foundation OperationsVMware Aria Operations but the check box to delete related objects is not selected in the **Delete All Accounts** dialog box.<br><br>An object deleted from an adapter data source is identified by VMware Cloud Foundation OperationsVMware Aria Operations as not existing and VMware Cloud Foundation OperationsVMware Aria Operations can no longer collect data about the object. Whether VMware Cloud Foundation OperationsVMware Aria Operations identifies deleted objects as not existing depends on the adapter. This feature is not implemented in some adapters.<br><br>For example, if the retention time is 360 hours and a virtual machine is deleted from a vCenter instance, the virtual machine |

*Table continued on next page*

*Continued from previous page*

| Setting | Default Value | Description |
|---|---|---|
| | | remains as an object in VMware Cloud Foundation OperationsVMware Aria Operations for 15 days before it is deleted.<br><br>This setting applies to objects deleted from the data source or server, not to any objects you delete from VMware Cloud Foundation OperationsVMware Aria Operations on the Inventory page.<br><br>You can define the number of hours per object type to retain objects that no longer exist and check for object type overrides. To add individual object types and set up their values, click the **Object Deletion Scheduling** icon. You can also edit or delete these object types. |
| Deletion Scheduling Interval | 24 hours | Determines the frequency to schedule deletion of resources. This setting works with the Deleted Objects setting to remove objects that no longer exist in the environment. VMware Cloud Foundation OperationsVMware Aria Operations transparently marks objects for removal that have not existed for the length of time specified under Deleted Objects. VMware Cloud Foundation OperationsVMware Aria Operations then removes the marked objects at the frequency specified under Deletion Scheduling Interval. |
| Object History | 90 days<br><br>60 days | Number of days to retain the history of the object configuration, relationship, and property data. |

*Table continued on next page*

*Continued from previous page*

| Setting | Default Value | Description |
|---|---|---|
| | | **NOTE**<br>The retention decision for each property value is based on the end date of that value instead of the start date. So any property value with an end date within the last 90 days interval (default setting) is retained. For example, if property value A was first published 1 year ago (which is long before the object history retention interval of 90 days) and property value B was published 3 days ago, and if no more points for that property exists, then the end date for the value A will be 3 days ago and for value B now. So both end dates will be within the recent 90 day limit and both values will be retained.<br><br>The retention decision for each property value is based on the end date of that value instead of the start date. So any property value with an end date within the last 60 days interval (default setting) is retained. For example, if property value A was first published 1 year ago (which is long before the object history retention interval of 60 days) and property value B was published 3 days ago, and if no more points for that property exists, then the end date for the value A will be 3 days ago and for value B now. So both end dates will be within the recent 60 day limit and both values will be retained.<br><br>The configuration data is the collected data from the monitored objects on which the metrics are based. The collected data includes changes to the configuration of the object.<br><br>The data is purged from the system after the specified number of days.<br><br>You can edit this setting for a value from 10 days to 60 days in VMware Cloud Foundation Operations. |
| Generated Reports Retention | deactivated<br><br>90 days | Number of days to retain generated reports. If deactivated, all the generated reports will be retained.<br><br>The minimum number of days you can set is one and the maximum number of days you can set is 3600.<br><br>Number of days to retain generated reports. |
| Symptoms/Alerts | 45 days<br><br>30 days | Number of days to retain canceled alerts and symptoms.<br><br>The alerts and symptoms are either canceled by the system or by a user.<br><br>You cannot edit this setting in VMware Cloud Foundation Operations. |

*Table continued on next page*

*Continued from previous page*

| Setting | Default Value | Description |
|---|---|---|
| Time Series Data Retention | 6 months<br><br>3 months | Number of months that you want to retain the collected and calculated metric data for the monitored objects. This setting is set to 6 months by default for 5 minutes interval data retention.<br><br>Number of months that you want to retain the collected and calculated metric data for the monitored objects. This setting is set to 3 months by default for 5 minutes interval data retention. |
| Additional Time Series Retention | 36 months<br><br>12 months | The number of months that the roll-up data extends beyond the regular period. The roll-up data is available starting from the end of the regular period and until the end of the roll-up data retention period. If you specify 0 as the value, then this will effectively deactivate the Additional Time Series Data Retention time and only data specified in Time Series Retention is stored. This setting ensures that after 6 months of normal retention for 5 minutes, the seventh month data is rolled up into a one hour roll up. You can set up this option up to 120 months for data roll ups.<br><br>The number of months that the roll-up data extends beyond the regular period. The roll-up data is available starting from the end of the regular period and until the end of the roll-up data retention period.<br><br>You can edit this setting for a value of 0 months to 12 months in VMware Cloud Foundation Operations. If you specify 0 as the value, then this will effectively deactivate the Additional Time Series Data Retention time and only data specified in Time Series Retention is stored. This setting ensures that after 3 months of normal retention for 5 minutes, the fourth month data is rolled up into a one hour roll up. |
| Near Real-Time Monitoring Data Retention | 3 days | The number of days to retain the near real-time data collected from the vCenter in VMware Cloud Foundation Operations. You cannot edit this setting in VMware Cloud Foundation Operations. |
| Deleted Users | 100 days | You can specify the number of days to keep custom content created by a user who has been removed from VMware Cloud Foundation OperationsVMware Aria Operations or by the automatic synchronization of LDAP. For example, the custom dashboards created by a user. |
| Configuration management drift history Data retention | 7 days | The number of days to retain historical data for the configuration drift report. |
| External Event Based Active Symptoms | deactivated | The number of days to retain the external event-based active symptoms. |

## Settings in the Cost/Price Category

**Table 241: Global Setting Default Values and Descriptions**

| Setting | Default Value | Description |
|---|---|---|
| Cost Calculation | activated | The host time at which cost calculation is run. |
| Currency | | You can specify the currency unit that is used for cost calculation. Click **Set Currency** and select the currency from the list of currency types. Click the check box to confirm your selection and then click **Set Currency**.<br><br>You can edit the currency to update the currency for cost calculation.<br><br>**NOTE**<br>Editing the currency has the following effects:<br><br>• Views, dashboards, and reports display data with the old and new currency and there will be discrepancies if you compare data across a timeline that includes currency change.<br>• Billing schedules get updated if the currency is changed between billing cycles. New scheduled bills are generated with the new currency and the currency change date as the start date. You must manually generate a bill for the timeline before that for the older currency.<br>• The cost and price metrics collected after the update display the new currency value. Older metrics that are already published continue to display the old currency value.<br>• The cost analysis report only displays data with the new currency if the analysis includes a time range with a currency change.<br>To view the currency change timeline, you can view the **Currency Code** property under the Universe object type. |
| Cluster Utilization Ceiling Factor | 5 | Ceiling for Expected Utilization when running on Actual Utilization. |
| Tag Based Costing Metrics | deactivated | When activated, VMware Cloud Foundation OperationsVMware Aria Operations additional cost metrics per tag. |
| Tag Based Pricing Metrics | deactivated | When activated, VMware Cloud Foundation OperationsVMware Aria Operations additional pricing metrics per tag. |

## Settings in the User Access Category

**Table 242: Global Setting Default Values and Descriptions**

| Setting | Entry | Description |
|---|---|---|
| Allow vCenter users to log in to individual vCenters from VMware Cloud Foundation OperationsVMware Aria Operations | activated | Allows vCenter users to log in to individual vCenters using the VMware Cloud Foundation OperationsVMware Aria Operations UI. |
| Allow vCenter users to log in from vCenter clients | activated | Allows vCenter users to log in from vCenter clients. |
| Allow vCenter users to log in to all vCenters using the VMware Cloud Foundation OperationsVMware Aria Operations UI | activated | Allows concurrent UI login sessions per user. Once changed, this setting affects the subsequent login sessions. |
| Concurrent UI login sessions | activated | Allows concurrent UI login sessions per user. Once changed, this setting affects the subsequent login sessions. |
| Allow non-imported vIDM user access | activated | Allows non-imported VMware Identity Manager users to be created automatically as read-only users upon first access. If deactivated, only VMware Identity Manager imported users or users belonging to imported VMware Identity Manager groups will be granted access. |
| Password Change | activated | Allows the users to modify the settings required to change their passwords. |
| Login Message | activated | Allows the user to display the login message when the user launches the login page in VMware Cloud Foundation OperationsVMware Aria Operations. You can set a message that requires explicit consent, for example terms and conditions or a Standard Mandatory DoD Notice. |

## Settings in the Reclamation Category

Settings in the reclamation category display information about powered -off VMs, idle VMs, snapshots, and orphaned disks.

**Table 243: Global Settings Default Values and Descriptions**

| Setting | Entry | Description |
|---|---|---|
| Include Powered off VMs for Reclamation | activated | VMs that have been continuously powered off during the defined time period.<br>The total storage capacity used is reclaimable. Total storage reclaimable cost is computed by multiplying |

*Table continued on next page*

*Continued from previous page*

| Setting | Entry | Description |
|---|---|---|
| | | storage rate with storage utilization. The direct cost of VM is also attributed. |
| Idle VM Management | | |
| Include Idle VMs for Reclamation | activated | You can configure the following parameters based on which VMware Aria OperationsVMware Cloud Foundation Operations calculates idle VMs:<br>• How many days the VM has been idle to be classified as a reclaimable idle VM.<br>• The number of days before which the provisioned VMs are not considered in the classification of reclaimable idle. |
| VM Idleness Criteria | 100 MHz of CPU for 100% of Time | Defines the criteria to identify VMs that remain idle during the defined time period.<br>You can set the following values:<br>• MHz of CPU consumption of VMs, and percentage of the time VMs have less than defined MHz of CPU within each day. |
| Include Snapshots for Reclamation | activated | VM snapshots that have existed for the entire defined time period.<br>Snapshots of a VM use storage space and such storage is reclaimable. The reclaimable cost is computed by multiplying the storage rate with the reclaimable storage value. |
| Orphaned Disk Management | | |
| Include Orphaned Disks for Reclamation | deactivated | Orphaned disks are VMDKs that are associated with a VM which are not in inventory, but still available in a data store. You can configure the minimum number of days for which VMDKs not related to any existing VM will be reported as orphaned and appear under the Orphaned Disks in the Reclaim page. |
| Orphaned Disks Collection Time | 8:00 pm | Host time to collect orphaned disks. |

*Table continued on next page*

*Continued from previous page*

| Setting | Entry | Description |
|---|---|---|
| | | **NOTE**<br>VMware Aria OperationsVMware Cloud Foundation Operations checks for orphaned VMDKs in the vSphere Client instances based on the set time. The **Orphaned Disks Collection** and the **Cost Calculation** settings are interrelated. It is recommended to schedule orphaned disks collection before cost calculation so that the cost is calculated based on the recent list of reclaimable orphaned VMDKs. The default value for Orphaned Disks Collection is set to 8:00 PM, and the default value for Cost Calculation is set to 9:00 PM.<br>You can navigate to the **Cost/Price** section under **Administration › Global Settings** and change the value of the **Cost Calculation** time. |

## Settings in the System Settings Category

**Table 244: Global Setting Default Values and Descriptions**

| Setting | Entry | Description |
|---|---|---|
| Session Timeout | 30 minutes<br><br>120 minutes | If your connection to VMware Cloud Foundation OperationsVMware Aria Operations is idle for the specified amount of time, you are logged out of the application.<br><br>You must provide credentials to log back in. |
| Dynamic Threshold Calculation | activated | Determines whether to calculate normal levels of threshold violation for all objects. |

*Table continued on next page*

*Continued from previous page*

| Setting | Entry | Description |
|---|---|---|
| | | If the setting is deactivated, the following area of VMware Cloud Foundation OperationsVMware Aria Operations does not work or are not displayed:<br>• Alert symptom definitions based on dynamic thresholds will not work<br>• Metric charts that display normal behavior are not present<br>Deactivate this setting only if you have no alternative options for managing resource constraints for your VMware Cloud Foundation OperationsVMware Aria Operations system. |
| Customer Experience Improvement Program | activated | Determines whether to participate in the Customer Experience Improvement Program by having VMware Cloud Foundation OperationsVMware Aria Operations send anonymous usage data to https://vmware.com.<br>**NOTE**<br>CEIP is activated by default when you install a new VMware Cloud Foundation OperationsVMware Aria Operations instance or when you upgrade to version 8.14 of VMware Cloud Foundation OperationsVMware Aria Operations. For more information, see 'Logging in to VMware Cloud Foundation OperationsVMware Aria Operations' in *Getting Started with VMware Cloud Foundation OperationsVMware Aria Operations Guide.*<br>Deactivate this setting only if you want to opt out of CEIP. |

*Table continued on next page*

*Continued from previous page*

| Setting | Entry | Description |
|---|---|---|
| | | **NOTE**<br>If you deactivate this setting in version 8.14 of VMware Cloud Foundation OperationsVMware Aria Operations, it stays deactivated for future releases. |
| System access URL | | You can specify the URL that is used to access the system when a load balancer is used. The URL that you enter here is displayed in the outbound notifications and while sharing dashboards. The IP /FQDN of the URL is used to register the vCenter and VMware Cloud Director when you configure the vCenter cloud account and VMware Cloud Director accordingly. |
| Docker image default registry URL | | You can specify the URL that is used to pull the docker images for containerized adapters. This URL can be overridden by installed management packs.<br>You can bring your own content using the containerized adapters. For more information, contact ops-integration-sdk-beta@vmware.com . |
| Automated Actions | activated | Determines whether to allow VMware Cloud Foundation OperationsVMware Aria Operations to automate actions. When an alert triggered, the alert provides recommendations for remediation. You can automate an action to remediate an alert when the recommendation is the first priority for that alert. You activate actionable alerts in your policies. |
| Credential ownership enforcement | activated | This setting enforces credential ownership and allows users to only access credentials they own.<br>Deactivate this option to modify credentials owned by other users. Once deactivated, you can view, or edit credentials created by other users, or delete credentials that are not in use. |

*Table continued on next page*

*Continued from previous page*

| Setting | Entry | Description |
|---|---|---|
| | | **NOTE**<br>After deactivating, if a user edits a credential owned by another user and then reactivates this option, the ownership of the credential transfers to the user who edited it last.<br><br>For more information on credentials, see Configuring Credentials in Integrations and Manage Credentials. |
| Activate Standard Certificate Validation | deactivated | This option activates certificate verification to Test Connection in the Create or Modify AI screen, using a standard verification flow.<br><br>The option checks CA authority.<br><br>• Certificate Subject DN<br>• Subject alternative name<br>• Certificate validity period<br>• Revocation list<br>This option also presents dialogs to user if one of those checks fail. It is up to the adapter implementation on how the adapter checks source certificate validity during a normal collection cycle. On a usual scenario, adapters just perform a thumb-print verification. However, in case this flag is activated, Test connection validates certificates in full scale and accepts certificates that are matching all criteria without any user dialogs. |
| Threshold For Adapters Certificate Expiration Alert | 5 for critical<br><br>14 for immediate<br><br>30 for warning | Set the number of days before which the system must raise certification expiration alerts.<br>To change the values of Critical, Immediate and Warning alerts, click the corresponding icon and move them along the slider. Alternatively, set the values manually.<br><br>Click **Save** after you make your changes. |

*Table continued on next page*

*Continued from previous page*

| Setting | Entry | Description |
|---|---|---|
| | | **NOTE**<br>For critical alerts, an alert banner is displayed under **Home › Overview** on the day of certification expiry. For example, if the critical alert is set to 10, then on the 10th day, an alert banner is displayed in the Home page. |

## Managing Users and Access Control in VMware Aria OperationsVMware Cloud Foundation Operations

Managing Users and Access Control
Managing Users and Access Control
To ensure security of the objects in your VMware Aria Operations instance, as a system administrator you can manage all aspects of user access control. You create user accounts, assign each user to be a member of one or more user groups, and assign roles to each user or user group to set their privileges.

To ensure security of the objects in your VMware Aria OperationsVMware Cloud Foundation Operations instance, as a system administrator you can manage some aspects of user access control. You can assign each user to be a member of one or more user groups, and assign roles to each user or user group to set their privileges.
Users must have privileges to access specific features in the VMware Aria OperationsVMware Cloud Foundation Operations user interface. Access control is defined by assigning privileges to both users and objects. You can assign one or more roles to users, and activate them to perform a range of different actions on the same types of objects. For example, you can assign a user with the privileges to delete a virtual machine, and assign the same user with read-only privileges for another virtual machine.

### User Access Control

You can authenticate users in VMware Aria OperationsVMware Cloud Foundation Operations in several ways.

In order to use your corporate credentials to log on to VMware Cloud Service including VMware Cloud Foundation Operations, you can federate Identity and Access Management to your corporate domain(s). For more information, see Setting Up Enterprise Federation with VMware Cloud Services in the VMware Cloud services Product Documentation.

- Create local user accounts in VMware Aria OperationsVMware Cloud Foundation Operations.
- Use VMware vCenter users. After the vCenter is registered with VMware Aria OperationsVMware Cloud Foundation Operations, configure the vCenter user options in the VMware Aria OperationsVMware Cloud Foundation Operations global settings to activate a vCenter user to log in to VMware Aria OperationsVMware Cloud Foundation Operations. When logged into VMware Aria OperationsVMware Cloud Foundation Operations, vCenter users access objects according to their vCenter-assigned permissions.
- Add an authentication source to authenticate imported users and user group information that resides on another machine.
  - Use LDAP to import users or user groups from an LDAP server. LDAP users can use their LDAP credentials to log in to VMware Aria OperationsVMware Cloud Foundation Operations.
  - Create a single sign-on source and import users and user groups from a single sign-on server. Single sign-on users can use their single sign-on credentials to log in to VMware Aria OperationsVMware Cloud Foundation Operations and vCenter. You can also use Active Directory through single sign-on by configuring the Active Directory through single sign-on and adding the single sign-on source to VMware Aria OperationsVMware Cloud Foundation Operations.

## User Preferences

To determine the display options for VMware Aria OperationsVMware Cloud Foundation Operations, such as colors for the display and health chart, the number of metrics and groups to display, and whether to synchronize system time with the host machine, you configure the user preferences on the top toolbar.

### Users of VMware Aria OperationsVMware Cloud Foundation Operations

Each user has an account to authenticate them when they log in to VMware Aria OperationsVMware Cloud Foundation Operations.

The accounts of local users and LDAP users are visible in the VMware Aria OperationsVMware Cloud Foundation Operations user interface when they are set up. The accounts of vCenter and single sign-on users only appear in the user interface after a user logs in for the first time. Each user can be assigned one or more roles, and can be an authenticated member of one or more user groups.

### Local Users in VMware Aria OperationsVMware Cloud Foundation Operations

Local Users
Local Users
When you create user accounts in a local VMware Aria OperationsVMware Cloud Foundation Operations instance, VMware Aria OperationsVMware Cloud Foundation Operations stores the credentials for those accounts in its global database, and authenticates the account user locally.

Each user account must have a unique identity, and can include any associated user preferences.

If you are logging in to VMware Aria OperationsVMware Cloud Foundation Operations as a local user, and on occasion receive an `invalid password` message, try the following workaround. In the Login page, change the Authentication Source to **All vCenter Servers**, change it back to **Local Users**, and log in again.

### vCenter Users in VMware Aria OperationsVMware Cloud Foundation Operations

vCenter Users
vCenter Users
VMware Aria OperationsVMware Cloud Foundation Operations supports vCenter users. To log in to VMware Aria OperationsVMware Cloud Foundation Operations, vCenter users must be valid users in vCenter.

## Roles and Associations

A vCenter user must have either the vCenter Admin role or one of the VMware Aria OperationsVMware Cloud Foundation Operations privileges, such as PowerUser which assigned at the root level in vCenter, to log in to VMware Aria OperationsVMware Cloud Foundation Operations. VMware Aria OperationsVMware Cloud Foundation Operations uses only the vCenter privileges, meaning the VMware Aria OperationsVMware Cloud Foundation Operations roles, at the root level, and applies them to all the objects to which the user has access. After logging in, vCenter users can view all the objects in VMware Aria OperationsVMware Cloud Foundation Operations that they can already view in vCenter.

## Logging in to vCenter Instances and Accessing Objects

vCenter users can access either a single vCenter instance or multiple vCenter instances, depending on the authentication source they select when they log in to VMware Aria OperationsVMware Cloud Foundation Operations.

- If users select a single vCenter instance as the authentication source, they have permission to access the objects in that vCenter instance. After the user has logged in, an account is created in VMware Aria OperationsVMware Cloud Foundation Operations with the specific vCenter instance serving as the authentication source.
- If users select **All vCenter Servers** as the authentication source, and they have identical credentials for each vCenter in the environment, they see all the objects in all the vCenter instances. Only users that have been authenticated by all the vCenter Servers in the environment can log in. After a user has logged in, an account is created in VMware Aria OperationsVMware Cloud Foundation Operations with all vCenter instances serving as the authentication source.

VMware Aria OperationsVMware Cloud Foundation Operations does not support linked vCenter instances. Instead, you must configure the vCenter adapter for each vCenter instance, and register each vCenter instance to VMware Aria OperationsVMware Cloud Foundation Operations.

Only objects from a specific vCenter instance appear in VMware Aria OperationsVMware Cloud Foundation Operations. If a vCenter instance has other linked vCenter instances, the data does not appear.

### vCenter Roles and Privileges

You cannot view or edit vCenter roles or privileges in VMware Aria OperationsVMware Cloud Foundation Operations. VMware Aria OperationsVMware Cloud Foundation Operations sends roles as privileges to vCenter as part of the vCenter Global privilege group. A vCenter administrator must assign VMware Aria OperationsVMware Cloud Foundation Operations roles to users in vCenter.

VMware Aria OperationsVMware Cloud Foundation Operations privileges in vCenter have the role appended to the name. For example, VMware Aria OperationsVMware Cloud Foundation Operations ContentAdmin Role, or VMware Aria OperationsVMware Cloud Foundation Operations PowerUser Role.

### Read-Only Principal

A vCenter user is a read-only principal in VMware Aria OperationsVMware Cloud Foundation Operations, which means that you cannot change the role, group, or objects associated with the role in VMware Aria OperationsVMware Cloud Foundation Operations. Instead, you must change them in the vCenter instance. The role applied to the root folder applies to all the objects in vCenter to which a user has privileges. VMware Aria OperationsVMware Cloud Foundation Operations does not apply individual roles on objects. For example, if a user has the PowerUser role to access the vCenter root folder, but has read-only access to a virtual machine, VMware Aria OperationsVMware Cloud Foundation Operations applies the PowerUser role to the user to access the virtual machine.

### Refreshing Permissions

When you change permissions for a vCenter user in vCenter, the user must log out and log back in to VMware Aria OperationsVMware Cloud Foundation Operations to refresh the permissions and view the updated results in VMware Aria OperationsVMware Cloud Foundation Operations. Alternatively, the user can wait for VMware Aria OperationsVMware Cloud Foundation Operations to refresh. The permissions refresh at fixed intervals, as defined in the `$ALIVE_BASE/user/conf/auth.properties` file. The default refreshing interval is half an hour. If necessary, you can change this interval for all nodes in the cluster.

### Single Sign-On and vCenter Users

When vCenter users log into VMware Aria OperationsVMware Cloud Foundation Operations by way of single sign-on, they are registered on the VMware Aria OperationsVMware Cloud Foundation Operations User Accounts page. If you delete the account of a vCenter user that has logged into VMware Aria OperationsVMware Cloud Foundation Operations by way of single sign-on, or remove the user from a single sign-on group, the user account entry still appears on the User Account page and you must delete it manually.

**Generating Reports**

vCenter users cannot create or schedule reports in VMware Aria OperationsVMware Cloud Foundation Operations.

### Backward Compatibility for vCenter Users in VMware Aria OperationsVMware Cloud Foundation Operations

Backward Compatibility for vCenter Users
Backward Compatibility for vCenter Users
VMware Aria OperationsVMware Cloud Foundation Operations provides backward compatibility for users of the earlier version of VMware Aria OperationsVMware Cloud Foundation Operations, so that users of vCenter who have privileges in the earlier version in vCenter can log in to VMware Aria OperationsVMware Cloud Foundation Operations.

When you register VMware Aria OperationsVMware Cloud Foundation Operations in vCenter, certain roles become available in vCenter.

- The Administrator account in the previous version of VMware Aria OperationsVMware Cloud Foundation Operations maps to the PowerUser role.
- The Operator account in the previous version of VMware Aria OperationsVMware Cloud Foundation Operations maps to the ReadOnly role.

During registration, all roles in VMware Aria OperationsVMware Cloud Foundation Operations, except for VMware Aria OperationsVMware Cloud Foundation Operations Administrator, Maintenance, and Migration, become available dynamically in vCenter. Administrators in vCenter have all of the roles in VMware Aria OperationsVMware Cloud Foundation Operations that map during registration, but these administrator accounts only receive a specific role on the root folder in vCenter if it is specially assigned.

Registration of VMware Aria OperationsVMware Cloud Foundation Operations with vCenter is optional. If users choose not to register VMware Aria OperationsVMware Cloud Foundation Operations with vCenter, a vCenter administrator can still use their user name and password to log in to VMware Aria OperationsVMware Cloud Foundation Operations, but these users cannot use the vCenter session ID to log in. In this case, typical vCenter users must have one or more VMware Aria OperationsVMware Cloud Foundation Operations roles to log in to VMware Aria OperationsVMware Cloud Foundation Operations.

When multiple instances of vCenter are added to VMware Aria OperationsVMware Cloud Foundation Operations, user credentials become valid for all of the vCenter instances. When a user logs in to VMware Aria OperationsVMware Cloud Foundation Operations, if the user selects all vCenter options during login, VMware Aria OperationsVMware Cloud Foundation Operations requires that the user's credentials are valid for all of the vCenter instances. If a user account is only valid for a single vCenter instance, that user can select the vCenter instance from the login drop-down menu to log in to VMware Aria OperationsVMware Cloud Foundation Operations.

vCenter users who log in to VMware Aria OperationsVMware Cloud Foundation Operations must have one or more of the following roles in vCenter:

- VMware Aria Operations Content Admin Role
- VMware Aria Operations General User Role 1
- VMware Aria Operations General User Role 2
- VMware Aria Operations General User Role 3
- VMware Aria Operations General User Role 4
- VMware Aria Operations Power User Role
- VMware Aria Operations Power User without Remediation Actions Role
- VMware Aria Operations Read Only Role

For more information about vCenter users, groups, and roles, see the vCenter documentation.

### External User Sources in VMware Aria OperationsVMware Cloud Foundation Operations

External User Sources
External User Sources
You can obtain user accounts from external sources so that you can use them in your VMware Aria OperationsVMware Cloud Foundation Operations instance.

There are two types of external user identity sources:
- Lightweight Directory Access Protocol (LDAP): Use the LDAP source if you want to use the Active Directory or LDAP servers as authentication sources. The LDAP source does not support multi-domains even when there is a two-way trust between Domain A and Domain B.
- Single Sign-On (SSO): Use a single sign-on source to perform single sign-on with any application that supports vCenter single sign-on, including VMware Aria OperationsVMware Cloud Foundation Operations. For example, you can install a standalone vCenter Platform Services Controller (PSC) and use it to communicate with an Active Directory server. Use a PSC if the Active Directory has a setup that is too complex for the simple LDAP source in VMware Aria OperationsVMware Cloud Foundation Operations, or if the LDAP source is experiencing slow performance.

### VMware Cloud Services Users

You can import and synchronize users and user groups of VMware Cloud Services Platform in vRealize Operations Cloud. You can assign roles and permissions to individual users of VMware Cloud Services Platform or you can assign the roles and permissions to a user group.

#### Operations Supported in vRealize Operations Cloud

You can perform the following operations after you log in to vRealize Operations Cloud as Administrator:
- View the users and user groups imported from VMware Cloud Services Platform under respective groups (Administration and User) in vRealize Operations Cloud.
- Assign roles and permission to the user and user groups that are automatically synchronized from VMware Cloud Services Platform.
- Ensure that users and user groups that are synchronized from VMware Cloud Services Platform have the same permissions in vRealize Operations Cloud.
- Change in user or user groups permission in VMware Cloud Services Platform is reflected in vRealize OperationsCloud.

#### Operations Not Supported in vRealize Operations Cloud

The following operations are not supported for users and user groups imported from VMware Cloud Services Platform:
- Modify permissions for users or user groups that are defined in VMware Cloud Services Platform.
- Delete users or user groups defined in VMware Cloud Services Platform.
- Move users from one user group to another.

### Roles and Privileges in VMware Aria OperationsVMware Cloud Foundation Operations

Roles and Privileges
Roles and Privileges
VMware Aria Operations provides several predefined roles to assign privileges to users. You can also create your own roles.

VMware Cloud Foundation Operations provides two predefined roles- the GeneralUser role and the Administrator role. These two roles are assigned to the user by the organization owner from VMware Cloud Services portal.
You must have privileges to access specific features in the VMware Cloud Foundation Operations user interface. The roles associated with your user account determine the features you can access and the actions you can perform.

Each predefined role includes a set of privileges for users to perform, create, read, update, or delete actions on components such as dashboards, reports, administration, capacity, policies, problems, symptoms, alerts, user account management, and adapters.

Each predefined role includes a set of privileges for users to perform

**Administrator**

Includes privileges to all features, objects, and actions in VMware Aria Operations.

Includes privileges to manage the VMware Cloud Foundation Operations instance and its objects. The Administrator can also customize the privileges associated with the GeneralUser role.

**PowerUser**

Users have privileges to perform the actions of the Administrator role except for privileges to user management and cluster management. VMware Aria OperationsVMware Cloud Foundation Operations maps vCenter users to this role.

**PowerUserMinusRemediation**

Users have privileges to perform the actions of the Administrator role except for privileges to user management, cluster management, and remediation actions.

**ContentAdmin**

Users can manage all content, including views, reports, dashboards, and custom groups in VMware Aria OperationsVMware Cloud Foundation Operations.

**GeneralUser-1 through GeneralUser-4**

**GeneralUser**

This role is defined role out of the box. This is the only default editable role, which has more permissions than ReadOnly user and less than Administrator.

These predefined template roles are initially defined as ReadOnly roles. vCenter administrators can configure these roles to create combinations of roles to give users multiple types of privileges. Roles are synchronized to vCenter once during registration.

**ReadOnly**

Users have read-only access and can perform read operations, but cannot perform write actions such as create, update, or delete.

### User Scenario: Manage User Access Control

As a system administrator or virtual infrastructure administrator, you manage user access control in VMware Aria OperationsVMware Cloud Foundation Operations so that you can ensure the security of your objects. Your company just hired a new person, and you must create a user account and assign a role to the account so that the new user has permission to access specific content and objects in VMware Aria OperationsVMware Cloud Foundation Operations.

As a system administrator or virtual infrastructure administrator, you manage user access control in VMware Aria OperationsVMware Cloud Foundation Operations so that you can ensure the security of your objects. Your company just hired a new person, and you must assign a role to the account so that the new user has permission to access specific content and objects in VMware Aria OperationsVMware Cloud Foundation Operations.

Verify that the following conditions are met:

- VMware Aria OperationsVMware Cloud Foundation Operations is installed and operating properly, and contains objects such as clusters, hosts, and virtual machines.
- One or more user groups are defined.

In this scenario you will learn how to create user accounts and roles, and assign roles to the user accounts to specify access privileges to views and objects. You will then demonstrate the intended behavior of the permissions on these accounts.

In this scenario you will learn how to assign roles to the user accounts to specify access privileges to views and objects. You will then demonstrate the intended behavior of the permissions on these accounts.
You will create a new user account, named Tom User, and a new role that grants administrative access to objects in the VMware Aria Operations Clusters. You will apply the new role to the user account.
Finally, you will import a user account from an external LDAP user database that resides on another machine to VMware Aria OperationsVMware Cloud Foundation Operations, and assign a role to the imported user account to configure the user's privileges.
Create a new role.

Assign a new role.

### Create a New Role

You use roles to manage access control for user accounts in VMware Aria OperationsVMware Cloud Foundation Operations.

Verify that you understand the context of this scenario. See User Scenario: Manage User Access Control. For information about roles and associated permissions, see KB 59484.

In this procedure, you will add a new role and assign administrative permissions to the role.

1. From the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Click the **Roles** tab.
3. Click the **Add** icon on the toolbar to create a role.
   The **Create Role** dialog box appears.
4. For the role name, type `admin_cluster`, then type a description and click **OK**.
   The admin_cluster role appears in the list of roles.
5. Click the **admin_cluster** role.
6. In the Details grid below, on the Permissions pane, click the **Edit** icon.
   The **Assign Permissions to Role** dialog box appears.
7. Select the **Administrative Access - all permissions** check box.
8. Click **Update**.
   This action gives this role administrative access to all the features in the environment.

Create a user account, and assign this role to the account.

### Create a User Account

As an administrator you assign a unique user account to each user so that they can use VMware Aria OperationsVMware Cloud Foundation Operations. While you set up the user account, you assign the privileges that determine what activities the user can perform in the environment, and upon what objects.

Create a new role. See Create a New Role.

In this procedure, you will create a user account, assign the `admin_cluster` role to the account, and associate the objects that the user can access while assigned this role. You will assign access to objects in the VMware Aria Operations Cluster. Then, you will test the user account to confirm that the user can access only the specified objects.

1. From the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Click the **User Accounts** tab.
3. Click the **Add** icon to create a new user account, and provide the information for this account.

| Option | Description |
|---|---|
| User Name | Type the user name to use to log in to VMware Aria OperationsVMware Cloud Foundation Operations. |
| Password | Type a password for the user. |
| Confirm Password | Type the password again to confirm it. |
| First Name | Type the user's first name. For this scenario, type `Tom`. |
| Last Name | Type the user's last name. For this scenario, type `User`. |
| Email Address | (Optional). Type the user's email address. |
| Description | (Optional). Type a description for this user. |
| Disable this user | Do not select this check box, because you want the user to be active for this scenario. |
| Require password change at next login | Do not select this check box, because you do not need to change the user's password for this scenario. |

4. Click **Next**.
   The list of user groups appears.
5. Select a user group to add the user account as a member of the group.
6. Click the **Objects** tab.
7. Select the **admin_cluster** role from the drop-down menu.
8. Select the **Assign this role to the user** check box.
9. In the Object Hierarchies list, select the **VMware Aria Operations Cluster** check box.
10. Click **Finish**.
    You created a new user account for a user who can access all the VMware Aria Operations Cluster objects. The new user now appears in the list of user accounts.
11. Log out of VMware Aria OperationsVMware Cloud Foundation Operations.
12. Log in to VMware Aria OperationsVMware Cloud Foundation Operations as Tom User, and verify that this user account can access all the objects in the VMware Aria OperationsVMware Cloud Foundation Operations Cluster hierarchy, but not other objects in the environment.
13. Log out of VMware Aria OperationsVMware Cloud Foundation Operations.

You used a specific role to assign permission to access all objects in the VMware Aria Operations Cluster to a user account named Tom User.

Import a user account from an external LDAP user database that resides on another machine, and assign permissions to the user account.

### Import a User Account and Assign Permissions

You can import user accounts from external sources, such as an LDAP database on another machine, or a single sign-on server, so that you can give permission to those users to access certain features and objects in VMware Aria OperationsVMware Cloud Foundation Operations.

- Configure an authorization source. See Authentication Sources .
- Configure an authorization source. See the VMware Aria OperationsVMware Cloud Foundation Operations Information Center.

1. Log out of VMware Aria OperationsVMware Cloud Foundation Operations, then log in as a system administrator.
2. From the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
3. On the toolbar, click the horizontal ellipsis and then click **Import from Source** icon.
4. Specify the options to import user accounts from an authorization source.

a) On the Import Users page, from the **Import From** drop-down menu, select an authentication source.
b) In the **Domain Name** drop-down menu, type the domain name from which you want to import users, and click **Search**.
c) Select the users you want to import, and click **Next**.
d) On the **Groups** tab, select the user group to which you want to add this user account.
e) Click the **Objects** tab, select the **admin_cluster** role, and select the **Assign this role to the user** check box.
f) In the Object Hierarchies list, select the **VMware Aria Operations Cluster** check box, and click **Finish**.

5. Log out of VMware Aria OperationsVMware Cloud Foundation Operations.
6. Log in to VMware Aria OperationsVMware Cloud Foundation Operations as the imported user.
7. Verify that the imported user can access only the objects in the VMware Aria Operations Cluster.

You imported a user account from an external user database or server to VMware Aria OperationsVMware Cloud Foundation Operations, and assigned a role and the objects the user can access while holding this role to the user.

You have finished this scenario.

**Configure a Single Sign-On Source in VMware Aria OperationsVMware Cloud Foundation Operations**

Configure a Single Sign-On Source
As a system administrator or virtual infrastructure administrator, you use single sign-on to activate SSO users to log in securely to your VMware Aria OperationsVMware Cloud Foundation Operations environment.

- Verify that the server system time of the single sign-on source and VMware Aria OperationsVMware Cloud Foundation Operations are synchronized. If you need to configure the Network Time Protocol (NTP), see Cluster and Node Maintenance.
- Verify that the server system time of the single sign-on source and VMware Aria OperationsVMware Cloud Foundation Operations are synchronized. If you need to configure the Network Time Protocol (NTP), see information about cluster and node maintenance in the *VMware Aria OperationsVMware Cloud Foundation Operations Getting Started Guide*.
- Verify that you have access to a Platform Services Controller through the vCenter. See the VMware vSphere Information Center for more details.

After the single sign-on source is configured, users are redirected to an SSO identity source for authentication. When logged in, users can access other vSphere components such as the vCenter without having to log in again.

1. Log in to VMware Aria OperationsVMware Cloud Foundation Operations as an administrator.
2. From the left menu, click **Administration › Control Panel**, and then click the **Authentication Sources** tile.
3. Click **Add**.
4. In the Add Source for User and Group Import dialog box, provide information for the single sign-on source.

| Option | Action |
|---|---|
| Source Display Name | Type a name for the import source. |
| Source Type | Verify that SSO SAML is displayed. |
| Host | Enter the IP address or FQDN of the host machine where the single sign-on server resides. If you enter the FQDN of the host machine, verify that every node in the VMware Aria OperationsVMware Cloud Foundation Operations cluster can resolve the single sign-on host FQDN. |

*Table continued on next page*

*Continued from previous page*

| Option | Action |
|---|---|
| Port | Set the port to the single sign-on server listening port. By default, the port is set to 443. |
| User Name | Enter the user name that can log into the SSO server. |
| Password | Enter the password. |
| Grant administrator role to VMware Cloud Foundation Operations for future configuration? | Select **Yes** so that the SSO source is reregistered automatically if you make changes to the VMware Aria OperationsVMware Cloud Foundation Operations setup. If you select **No**, and the VMware Aria OperationsVMware Cloud Foundation Operations setup is changed, single sign-on users will not be able to log in until you manually reregister the single sign-on source. |
| Automatically redirect to vRealize Operations single sign-on URL? | Select **Yes** to direct users to the vCenter single-sign on log in page. If you select **No**, users are not redirected to SSO for authentication. |
| Import single sign-on user groups after adding the current source? | Select **Yes** so that the wizard directs you to the Import User Groups page when you have completed the SSO source setup. If you want to import user accounts, or user groups at a later stage, select **No**. |
| Advanced options | If your environment uses a load balancer, enter the IP address of the load balancer. |

5. Click **Test** to test the source connection, and then click **OK.**
   The certificate details are displayed.
6. Select the **Accept this Certificate** check box, and click **OK**.
7. In the Import User Groups dialog box, import user accounts from an SSO server on another machine.

| Option | Action |
|---|---|
| Import From | Select the single sign-on server you specified when you configured the single sign-on source. |
| Domain Name | Select the domain name from which you want to import user groups. If Active Directory is configured as the LDAP source in the PSC, you can only import universal groups and domain local groups if the vCenter resides in the same domain. |
| Result Limit | Enter the number of results that are displayed when the search is conducted. |
| Search Prefix | Enter a prefix to use when searching for user groups. |

8. In the list of user groups displayed, select at least one user group, and click **Next**.
9. In the Roles and Objects pane, select a role from the **Select Role** drop-down menu, and select the **Assign this role to the group** check box.
10. Select the objects users of the group can access when holding this role.

    To assign permissions so that users can access all the objects in VMware Aria OperationsVMware Cloud Foundation Operations, select the **Allow access to all objects in the system** check box.

11. Click **OK.**
12. Familiarize yourself with single-sign on and confirm that you have configured the single sign-on source correctly.
    a) Log out of VMware Aria OperationsVMware Cloud Foundation Operations.

b) Log in to the vSphere Web Client as one of the users in the user group you imported from the single sign-on server.
c) In a new browser tab, enter the IP address of your VMware Aria OperationsVMware Cloud Foundation Operations environment.
d) If the single sign-on server is configured correctly, you are logged in to VMware Aria OperationsVMware Cloud Foundation Operations without having to enter your user credentials.

**Edit a Single Sign-On Source**

Edit a single sign-on source if you need to change the administrator credentials used to manage the single sign-on source, or if you have changed the host of the source.

When you configure an SSO source, you specify either the IP address or the FQDN of the host machine where the single sign-on server resides. If you want to configure a new host, that is, if the single sign-on server resides on a different host machine than the one configured when the source was set up, VMware Aria Operations removes the current SSO source, and creates a new source. In this case, you must reimport the users you want to associate with the new SSO source.
If you want to change the way the current host is identified in VMware Aria Operations, for example, change the IP address to the FQDN and the reverse, or update the IP address of the PSC if the IP address of the configured PSC has changed, VMware Aria Operations updates the current SSO source, and you are not required to reimport users.

1. Log in to VMware Aria Operations as an administrator.
2. From the left menu, click **Administration › Control Panel**, and then click the **Authentication Sources** tile.
3. Select the single sign-on source, click the vertical ellipsis and then click **Edit**.
4. Make changes to the single sign-on source, and click **OK**.
   If you are configuring a new host, the New Single Sign-On Source Detected dialog box appears.
5. Enter the administrator credentials that were used to set up the single sign-on source, and click **OK**.
   The current SSO source is removed, and a new one created.
6. Click **OK** to accept the certificate.
7. Import the users you want to associate with the SSO source.

**Access Control in VMware Aria OperationsVMware Cloud Foundation Operations**

Access Control
Access Control
Each user must have a unique account with one or more roles assigned to enforce a role-based security when they use VMware Aria Operations. You create a user account, and assign the account to be a member of one or more user groups to allow the user to inherit the roles and scopes associated with the user group.

Each user must have a unique account with one or more roles assigned to enforce a role-based security when they use VMware Cloud Foundation Operations. Contact your organization owner to create user accounts for VMware Cloud Foundation Operations. Once created, you can assign the account to be a member of one or more user groups to allow the user to inherit the roles and scopes associated with the user group.

**Where You Find the Access Control Options**

You can manage user accounts and their associated user groups, roles, scopes and passwords.

From the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.

**Table 245: Access Control Tabs**

| Option | Description |
|---|---|
| User Accounts | Add, edit, remove, or import VMware Aria Operations user accounts from an LDAP database, and manage user roles, their membership in groups, and the scopes assigned for association with the user. Import user accounts from an LDAP database that resides on another machine.<br><br>Edit VMware Cloud Foundation Operations user accounts and manage user roles, their membership in groups, and the scopes assigned for association with the user.<br><br>vCenter users who are logged in to VMware Aria OperationsVMware Cloud Foundation Operations, either logged in directly or through the vSphere Client, appear in the list of user accounts. |
| User Groups | Add, edit, or remove, or import VMware Aria Operations user groups, update the members in a group and the associated scopes that they can access. Import user groups from an LDAP database or a single sign-on database that resides on another machine.<br><br>Add, edit, delete, or clone VMware Cloud Foundation Operations user groups, update the members in a group and the associated roles and scopes that they can access.<br><br>VMware Aria OperationsVMware Cloud Foundation Operations continuously synchronizes the user membership of imported LDAP user groups when the autosync option is enabled in the LDAP configuration. |
| Roles | For users to perform actions in VMware Aria OperationsVMware Cloud Foundation Operations, they must be assigned specific roles. With role-based access, when you assign a role to a user, you are determining not only what actions the user can perform in the system, but also the objects upon which those actions can be performed while holding the role. For example, to import or export a policy, the role assigned to your user account must have the Import or Export permissions enabled for policy management. |
| Scopes | Add, edit, clone, or remove scope associated with users or groups in VMware Aria OperationsVMware Cloud Foundation Operations. Scope lets you limit the access of a user or a set of users to VMware Aria OperationsVMware Cloud Foundation Operations. You can also define the scope for all the objects managed by VMware Aria OperationsVMware Cloud Foundation Operations. |

### Access Control: User Accounts Tab

You can add, edit, or remove VMware Aria Operations user accounts, and import user accounts from an external LDAP database. With access control, you manage roles, the objects a user can access while assigned a specific role, and the membership in user groups.

With access control, you manage roles, the objects a user can access while assigned a specific role, and the membership in user groups.

#### Where You Manage User Accounts

From the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.

To view the details of the user account configured in VMware Aria OperationsVMware Cloud Foundation Operations, click the user account, the user account details is displayed in the right-side panel. The user account details include the user groups associated with the user account along with the roles and scope of the selected user account.

**NOTE**
For user accounts only the exclusive roles that is associated with the user account is displayed. The roles inherited from the user groups are not displayed.

**Table 246: Access Control User Accounts Summary Grid**

| Summary Grid Options | Description |
|---|---|
| User Accounts toolbar | To manage user accounts, use the toolbar options.<br>• Click the **Add** button to add a user account, and provide the details for the user account in the **Account Information** page.<br>• Click the **Vertical Ellipses** to perform any one of the following actions:<br>– **Edit**. Edit the selected user account, and modify the details for the user account in the **Edit User Account** dialog box.<br>– **Export**. Click **Export User Accounts** to export a user account, and provide the details to export the user account in the **Export Authentication Source** dialog box.<br>– **Delete**. Delete a user account.<br>• Click the **Horizontal Ellipses** and select **Delete**. Delete a user account.<br>– **Export**. Click export to export a user account, and provide the details to export the user account in the **Export Authentication Source** dialog box.<br>– Click **Import** to import a user account, and provide the details to import the user account in the **Import User Accounts** dialog box.<br>– Click **Import from source** to import a user account, and provide the details to import the user group in the **Import User** dialog box. |
| User Name | User name, without spaces, that you use to log in to VMware Aria OperationsVMware Cloud Foundation Operations |
| First Name | User's first name, created when you create the user account. |
| Last Name | User's last name, created when you create the user account. |
| Email | User's email address, created when you create the user account. |
| Description | Description of the user account, defined when you create the user account. This information can identify the type of user and a summary of their access privileges. |
| Imported | Indicates whether the user account is imported or not. |
| Source Type | Indicates whether the user account is a local user, or an external user who is integrated through an external authentication source, such as from LDAP, SSO, AD, OpenLDAP, vCenter.<br><br>Indicates whether the user account platform. |
| Activated | Indicates whether the user account is activated to use VMware Aria OperationsVMwar e Cloud Foundation Operations features. An administrator can edit a user account to manually activate it, or deactivate it to prevent user access to VMware Aria OperationsV Mware Cloud Foundation Operations. |
| Locked | Indicates whether VMware Aria OperationsVMware Cloud Foundation Operations has locked the user account. For example, a user account can get locked based on the password lockout policy, or if the user enters an incorrect password three times in the span of five minutes. |
| Access All Objects | Indicates whether the user account is allowed to access all the objects that are imported into the VMware Aria OperationsVMware Cloud Foundation Operations instance. |

*Table continued on next page*

*Continued from previous page*

| Summary Grid Options | Description |
|---|---|
| Modified By | Indicates the last person to update the user account. |
| Last Modified | Indicates the last time the user account was updated. |
| Last Login Time | Indicates the last time the user logged in. |
| Filters | Activates you to search the list of user accounts according to the following criteria:<br>• User Name<br>• First Name<br>• Last Name<br>• Description<br>• Activated<br>• Locked<br>• Access All Objects<br>• Modified By<br>• Last Login Time |

After you add a user account, use the Details grid to view and edit which user accounts are assigned to user groups, and view the permissions assigned to the user account.

Use the Details grid to view and edit which user accounts are assigned to user groups, and view the permissions assigned to the user account.

**Table 247: Access Control User Accounts Details Grid**

| Details Grid Options | Description |
|---|---|
| First Name | User's first name, created when you create the user account. |
| Last Name | User's last name, created when you create the user account. |
| Modified By | Displays the name of the user who last updated the user account. |
| Last Modified | Displays the date and time when the user account was last updated. |
| Roles and Scope | Displays the role and scope associated with the selected user.<br>• Roles: Indicates the name of the role or roles assigned to the user.<br>• Scope : Indicates the scope associated with the user account. |
| User Groups | Assigned user groups appear when you click a user in the summary grid. You can then view and modify which user groups the user is associated with. |

## Modify User Accounts and Assign Groups and Permissions

You can add user accounts so that users can access the features of VMware Aria Operations and certain objects in the environment. Or, modify user accounts to change their attributes, deactivate or lock the accounts, or require them to change their password. After you add user accounts, you can assign them to one or more user groups, and assign roles and scopes. Assign the administrators role only to specific users who must access objects and perform actions in the entire environment.

Contact your organization owners to add user accounts so that users can access the features of VMware Aria OperationsVMware Cloud Foundation Operations and certain objects in the environment. After the organization owners add user accounts, you can assign them to one or more user groups, and assign roles and scopes. Assign the administrators role only to specific users who must access objects and perform actions in the entire environment.
Where You Add or Edit User Accounts

1. To add a user account, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. To modify a user account, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
3. In the **User Accounts** tab, click **Add**.
4. To edit a user account, click the vertical ellipsis and select **Edit**. You can also click the **EDIT ACCOUNT** button in the **User Details** page and edit the user account.
   The **Edit User Information** page is displayed.

   - By default, the user is added to the "Everyone" group. You can add or remove users to or from the user group when you create or edit the user group, . You can assign the permissions to the users by selecting the required role-scope pairs.
   - When you add users to a user group, the user inherits the permissions (role-scope pairs) from that group. You also have the option of not assigning permissions to the user, in that case the user actions are restricted, and the user cannot login to VMware Aria OperationsVMware Cloud Foundation Operations.
   - You cannot change the user name once you create it. So, select your username appropriately.
   - You have read only access for the Super Administrator account. You can change the password for the Super Administrator account from the user interface only.

**Table 248: Add or Edit Users Accounts- User Details Page**

| User Details Options | Description |
|---|---|
| User Name | User name, without spaces to access the VMware Aria OperationsVMware Cloud Foundation Operations |
| Password | User's password to access the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Confirm Password | Confirmation of the user's password. |
| Require password change at next login | Users can change their password the next time they log in to the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Deactivate this user | Deactivate the user account so that a user cannot access the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| First Name | User's first name, created when you create the user account. |
| Last Name | User's last name, created when you create the user account. |
| Email Address | User's email address, created when you create the user account. |
| Description | Description of the user account, defined when you create the user account. This information can identify the type of user and a summary of their access rights. |

**Assign Roles and Scope**

The Assign Roles and Scope option lets you select a role for each user and then assign a scope for that role.

5. From the **Select Role** drop-down menu, select a role for the user.
6. From the **Select Scope** drop-down menu, assign a scope for the selected user.

   **NOTE**
   You can click the + sign to add multiple roles and then assign the required scope for each role.

**Assign User Groups**

The Assign User Group option lets you select user groups that the user will be a member of.

7. Select the user groups for which you want the user to be member of.
8. Click **Save**.

   **NOTE**
   You cannot manually add users to groups imported from LDAP and SSO.

## Import User Accounts From Source

You can import user accounts so that users can access the features of VMware Aria OperationsVMware Cloud Foundation Operations and the objects in the environment. After you import user accounts, you can add them to non-imported user groups, assign roles and scopes to them.

Where You Import User Accounts

1. To import user accounts, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Click the horizontal ellipsis next to **Add** and then, click **Import from Source** to import user from another source.

**Table 249: Import Users from a LDAP Source**

| User Details Options | Description |
|---|---|
| Import From | Select the authentication source from the drop-down menu, the options are LDAP host machine, VMware Identity Manager, Active Directory, or Other sources configured to import user accounts. To know more about authentication,see Authentication Sources. |
| User Name | Click **Use alternate credentials** to display the user name of the LDAP source credential used to import user accounts to the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Password | Password for the LDAP source credential to import user accounts to the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Basic | Use the basic import setting with search option to look for user accounts. |
| Advanced | Displays the advanced import settings.<br>• Group Search Criteria. Search criteria to find LDAP groups. If not included, VMware Cloud Foundation OperationsVMware Aria Operations uses the default search parameters: `(|(objectClass=group)(objectClass=groupOfNames))` |

*Table continued on next page*

*Continued from previous page*

| User Details Options | Description |
|---|---|
| | • Member Attribute. Name of the attribute for a group object that contains the list of members. If not included, VMware Cloud Foundation OperationsVMware Aria Operations uses member by default.<br>• User Search Criteria. Search criteria to use the member field to find and cache LDAP users. You enter sets of key=value pairs in the form `(\|(key1=value1)(key2=value2))`. If not included, VMware Cloud Foundation OperationsVMware Aria Operations searches for each user separately. This operation might take extra time.<br>• Member Match Field. Name of the attribute for a user object to match with the member entry from a group object. If not included, VMware Cloud Foundation OperationsVMware Aria Operations treats the member entry as a distinguished name.<br>• LDAP Context Attributes. Attributes that VMware Cloud Foundation OperationsVMware Aria Operations applies to the LDAP context environment. You enter sets of key=value pairs separated by commas, such as **`java.naming.referral=ignore,java.naming.ldap.deleteRDNfalse`**. |
| Search String | Enter a search string, and click **Search** to start the search for user accounts. |
| User Name Summary grid | Lists the users available for import. Select the check box for each user to import, or select the **User Name** check box to import all users. User accounts that are already imported to VMware Aria OperationsVMware Cloud Foundation Operations do not appear in the list. |

**Table 250: Import Users from a VMware SSO Source**

| User Details Options | Description |
|---|---|
| Import From | VMware SSO is configured as the source to import user accounts. |
| Search Prefix | Enter a search string, and click **Search** to start the search for user accounts. |
| User Name Summary grid | Lists the users available for import. Select the check box for each user to import, or select the **User Name** check box to import all users. To appear in the list, the user configuration must be set to primary group in the default domain user group. User accounts that are already imported to VMware Aria OperationsVMware Cloud Foundation Operations do not appear in the list.<br>**NOTE**<br>After importing the user, you must edit the user information and assign a role and a scope. |

**Table 251: Import Users from a VMware Identity Manager Source**

| User Details Options | Description |
|---|---|
| Import From | VMware Identity Manager configured as the source to import user accounts. |

*Table continued on next page*

*Continued from previous page*

| User Details Options | Description |
|---|---|
| Search Option | Enter the search option, it can be a domain or search prefix. |
| Domain Name | Enter the domain name for import. |
| Search Prefix | Enter a search string, and click **Search** to start the search for user accounts. |
| User Name Summary grid | Lists the users available for import. Select the check box for each user to import, or select the **User Name** check box to import all users. To appear in the list, the user configuration must be set to primary group in the default domain user group. User accounts that are already imported to VMware Aria OperationsVMw are Cloud Foundation Operations do not appear in the list. |

**Table 252: Import Users from a Single Sign On Source**

| User Details Options | Description |
|---|---|
| Import From | SSO source configured as the source to import user accounts. |
| Domain Name | Select the domain name for import from the drop-down menu. |
| Result Limit | Determines the number of users displayed. |
| Search Prefix | Enter a search prefix, and click **Search** to start the search for user accounts. |
| User Name Summary grid | Lists the users available for import. Select the check box for each user to import, or select the **User Name** check box to import all users. To appear in the list, the user configuration must be set to primary group in the default domain user group. User accounts that are already imported to VMware Aria OperationsVMw are Cloud Foundation Operations do not appear in the list. |

3. After you enter the import users details, click **Next**

**Table 253: Import Users Accounts- Assign Groups Page**

| Assign Groups | Description |
|---|---|
| Assign Groups | Select or deselect the groups associated with the user account. To select or deselect all accounts, click the **Group Name** check box. You cannot add user accounts to groups imported from LDAP. |

4. Click **Finish**.

### Export and Import of User Accounts

You can export User account configurations from one VMware Aria OperationsVMware Cloud Foundation Operations and import into any VMware Aria OperationsVMware Cloud Foundation Operations.

Export User Accounts

1. From the left menu click **Administration › Control Panel**, and then click **Access Control** tile.
2. In the **User Accounts** tab, select the user accounts to be exported. Click the horizontal ellipsis next to **Add** and then, click **Export**.
3. You will be prompted to enter a password when you export the user accounts. Enter the password and note it down, you have to use the same password when you import the user accounts.
4. Click **Export**.

The user accounts `.json` file is downloaded to the default download location.

Import User Accounts

1. From the left menu click **Administration › Control Panel**, and then click **Access Control** tile.
2. In the **User Accounts** tab, click the horizontal ellipsis next to **Add** and then, click **Import**.
3. Click **Browse** and select the user accounts `.json` file.
4. Enter the same password which you had used during user accounts export.
5. In case of a conflict, select either **Overwrite existing User Accounts** or **Skip User Accounts**.
6. Click **Import**.

Important Points

- User scopes will be matched in the target VMware Aria OperationsVMware Cloud Foundation Operations, which means the user scopes references are being exported, then assigned to user while importing only if those scopes exist on the target VMware Aria OperationsVMware Cloud Foundation Operations.
- User roles will be matched in the target VMware Aria OperationsVMware Cloud Foundation Operations, which means user roles references are being exported, then assigned to user while importing only if those roles exist on the target VMware Aria OperationsVMware Cloud Foundation Operations.
- User groups will be matched in the target VMware Aria OperationsVMware Cloud Foundation Operations, which means user group references are being exported, then assigned to user while importing only if those user groups exist on the target VMware Aria OperationsVMware Cloud Foundation Operations.
- Import of the external users fails if their authentication source is not found in the target VMware Aria OperationsVMware Cloud Foundation Operations.
- The traversal specifications are not imported or exported in VMware Aria OperationsVMware Cloud Foundation Operations.
- An error message is displayed with details for failed imports.
- User Account Export or Import is not supported in VMware Aria Operations (SaaS).

### Access Control: User Groups Tab

You can manage the user groups associated with the users and objects in your environment. You can import user groups from an LDAP database that resides on another machine, or from a single sign-on server or from a source in VMware Identity Manager.

You can manage the user groups associated with the users and objects in your environment.

#### Where You Manage User Groups

1. To manage user groups, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Click the **User Groups** tab.

To view the details of the user groups configured in VMware Aria OperationsVMware Cloud Foundation Operations, click the user group, the user group details are displayed in the right-side panel. The user group details include the user accounts associated with the selected user group along with the roles and scope associated with the user group.

**Table 254: Access Control User Groups Summary Grid**

| Option | Description |
|---|---|
| User Groups toolbar | To manage user groups, use the toolbar options. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Click the **Add** button to add a user group, and provide the details for the user group in the **Group Information** page.<br>• Click the **Vertical Ellipses** to perform any one of the following actions:<br>– **Edit**. Edit the selected user group, and modify the details for the user group in the **Edit User Group** dialog box.<br>– **Delete**. Delete a user group.<br>– **Clone**. Clone a user group, and enter a name and description for the cloned user group.<br>– **Export**. Click export to export a user group, and provide the details to export the user group in the **Export Authentication Source** dialog box.<br>• Click the **Horizontal Ellipses** and select the available options.<br>– **Delete**. Click delete to Delete a user group.<br>– **Export**. Click export to export a user group, and provide the details to export the user group in the **Export Authentication Source** dialog box.<br>– Click **Import** to import a user group, and provide the details to import the user group in the **Import User Groups** dialog box.<br>– Click **Import from source** to import a user group, and provide the details to import the user group in the **Import User** dialog box. |
| Group Name | Name of the user group. |
| Description | Description of the group, indicating its purpose. |
| Members | Number of members in the group. |
| Group Type | Type of group, either a local user group or a group imported from LDAP.<br><br>Type of group like a local user group |
| Distinguished Name | Names for LDAP objects, such as domains and users.<br><br>Names of users. |
| Access All Objects | Indicates if the user group account is allowed to access all the objects that are imported into the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Modified By | Indicates the last person to update the user group. |
| Last Modified | Indicates the last time the user group was updated. |
| Filters | Enables you to search the list of user groups according to the following criteria:<br>• Group Name<br>• Description<br>• Access All Objects<br>• Modified By |

After you select a user group in the summary grid, view details about associated users in the Details pane.

**Table 255: Access Control User Groups Details Grid**

| Option | Description |
|---|---|
| User Group Name | User Group name, created when you create the user group. |
| Modified By | Displays the name of the user who last updated the user group. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Last Modified | Displays the date and time when the user group was last updated. |
| Description | Description of the group, indicating its purpose. |
| Roles and Scope | Displays the role and scope associated with the selected user.<br>• Role: Indicates the name of the role or roles assigned to the user.<br>• Scope: Indicates the scope associated with the user group. |
| User Accounts | Displays the user accounts associated with the selected user group. |

### Add User Groups and Assign Members and Roles

You can view and modify the details for user groups, including users, roles, and scope.

Where You Add User Groups

1. To add a user group, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Select the **User Groups** tab and then click **Add**
3. To edit a user group, click the vertical ellipsis and select **Edit**. You can also click the **EDIT GROUP** button in the **Group Details** page and edit the user group.

   **Table 256: Add or Edit Group Information - Name and Description Page**

   | Option | Description |
   |---|---|
   | Group Name | Name of the user group, either created manually, imported from a single sign-on server, or imported from an LDAP database that resides on another machine.<br>Name of the user group that is created manually. |
   | Description | Description of the user group, indicating its purpose. |

4. **Assign Roles and Scope**
   The Assign Roles and Scope option lets you select a role for the user group and then assign a scope for each role.
5. From the **Select Role** drop-down menu, select a role for the user group.
6. From the **Select Scope** drop-down menu, assign scope for the selected user group.

   **NOTE**
   You can click the + sign to add multiple roles and then assign the required scope for each role.

   **Assign Users**

   The Assign Users option lets you add members to the group based on the membership criteria that you have defined.
7. Select the users who would be the members of this group.
8. Click **Save**.

### Import User Groups From Source

You import user groups from a single sign-on server, VMware Identity Manager, Active Directory, or an LDAP database on another machine so that you can use those groups in VMware Aria OperationsVMware Cloud Foundation Operations.

Where You Import User Groups

1. To import a user group, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Select **User Groups** tab, click the horizontal ellipsis next to **ADD** button and select **Import**

   **NOTE**
   You can edit the imported group to assign Roles and Scopes.

   The options displayed in the Import User Groups page depend upon the authentication source you select.

**Table 257: Import User Groups Page - LDAP, Active Directory, and Others Sources**

| Option | Description |
|---|---|
| Import From | Host machine configured as the source to import the user groups. These options are displayed when the host machine of an LDAP, Active Directory, or Other source is selected. |
| User Name | User name of the source credential to import user groups to the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Password | Password for the source credential to import user groups to the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Search String | Invoke the search for user groups. |
| Basic | Use the basic import setting with search option to look for user groups. |
| Advanced | Displays the advanced import settings.<br>• Group Search Criteria. Search criteria to find LDAP groups. If not included, VMware Cloud Foundation OperationsVMware Aria Operations uses the default search parameters: `(|(objectClass=group)(objectClass=groupOfNames))`<br>• Member Attribute. Name of the attribute for a group object that contains the list of members. If not included, VMware Cloud Foundation OperationsVMware Aria Operations uses member by default.<br>• User Search Criteria. Search criteria to use the member field to find and cache LDAP users. You enter sets of key=value pairs in the form `(|(key1=value1)(key2=value2))`. If not included, VMware Cloud Foundation OperationsVMware Aria Operations searches for each user separately. This operation might take extra time.<br>• Member Match Field. Name of the attribute for a user object to match with the member entry from a group object. If not included, VMware Cloud Foundation OperationsVMware Aria Operations treats the member entry as a distinguished name.<br>• LDAP Context Attributes. Attributes that VMware Cloud Foundation OperationsVMware Aria Operations applies to the LDAP context environment. You enter sets of key=value pairs separated by commas, such as **`java.naming.referral=ignore,java.naming.ldap.deleteRDNfalse`**. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Group Name | Displays a list of user groups. Select the **Group Name** check box to import all the displayed user groups, or select the check box next to each user group that you want to import. |

**Table 258: Import User Groups Page - Single Sign On Source**

| Option | Description |
|---|---|
| Import From | Host machine configured as the source to import the user groups. |
| Domain Name | User name of the source credential to import user groups to the VMware Aria OperationsVMware Cloud Foundation Operations instance. |
| Result Limit | Determines the number of groups displayed. |
| Search Prefix | Enter a search prefix to narrow your search. |
| Group Name | Displays a list of user groups. Select the **Group Name** check box to import all the displayed user groups, or select the check box next to each user group that you want to import. |

**Table 259: Import User Groups from a VMware SSO Source**

| User Detail Options | Description |
|---|---|
| Import From | VMware SSO is configured as the source to import user groups. |
| Search Prefix | Enter a search string, and click **Search** to start the search for user groups. |
| Group Name Summary grid | Lists the users available for import. Select the check box for each user group to import, or select the **Group Name** check box to import all groups. User groups that are already imported to VMware Aria OperationsVMware Cloud Foundation Operations do not appear in the list.<br>**NOTE**<br>After importing the user group, you must edit the user group information and assign a role and a scope. |

**Table 260: Import User Groups from a VMware Identity Manager Source**

| User Details Options | Description |
|---|---|
| Import From | VMware Identity Manager configured as the source to import user groups. |
| Domain Name | Enter the domain name for import. |
| Search Prefix | Enter a search string, and click **Search** to start the search for user groups. |
| Group Name Summary grid | Lists the users available for import. Select the check box for each user group to import, or select the **Group Name** check box to import all groups. User groups that are already imported to VMware Aria OperationsVMware Cloud Foundation Operations do not appear in the list. |

3. After you enter the import user group details, click **Next**.
4. Click **Finish**.

### Export and Import of User Groups

You can export User Group configurations from one VMware Aria OperationsVMware Cloud Foundation Operations and import into any VMware Aria OperationsVMware Cloud Foundation Operations.

Export User Groups

1. From the left menu click **Administration › Control Panel**, and then click **Access Control** tile.
2. In the **User Groups** tab, select the user groups to be exported. Click the horizontal ellipsis next to **Add** and then, click **Export**.

The user groups `.json` file is downloaded to the default download location.

Import User Groups

1. From the left menu click **Administration › Control Panel**, and then click **Access Control** tile.
2. In the **User Groups** tab, click the horizontal ellipsis next to **Add** and then, click **Import**.
3. Click **Browse** and select the user groups `.json` file.
4. In case of a conflict, select either **Overwrite existing User Groups** or **Skip User Groups**.
5. Click **Import**.

Important Points

- User group’s scopes are matched in the target VMware Aria OperationsVMware Cloud Foundation Operations, which means user group’s scopes references are exported, then assigned to the user group while importing only if those scopes exist on the target VMware Aria OperationsVMware Cloud Foundation Operations.
- User group's roles are matched in the target VMware Aria OperationsVMware Cloud Foundation Operations, which means user group's role references are exported, then assigned to the user group while importing only if those roles exist on the target VMware Aria OperationsVMware Cloud Foundation Operations.
- User group's members are matched in the target VMware Aria OperationsVMware Cloud Foundation Operations, which means user group's member references are exported, then assigned to the user group while importing only if those users exist on the target VMware Aria OperationsVMware Cloud Foundation Operations.
- Import of the external user groups fails if their authentication source is not found in the target VMware Aria OperationsVMware Cloud Foundation Operations.
- The traversal specifications are not imported or exported in VMware Aria OperationsVMware Cloud Foundation Operations.
- An error message is displayed with details for failed imports.
- User groups of type “Cloud Services Platform” can not be exported in VMware Aria Operations (SaaS).

### Access Control: Roles Tab

You can assign users-specific roles to perform actions and view features and objects in VMware Cloud Foundation OperationsVMware Aria Operations. With role-based access, users can only perform the actions that their permissions allow.

#### Where You Manage User Roles

1. To manage user roles, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.

2. Click the **Roles** tab.

To view the details of roles configured in VMware Aria OperationsVMware Cloud Foundation Operations, click the role, the role details are displayed in the right-side panel. The role details include the permissions, user accounts, and user groups associated with the selected role.

You can view and edit details about a role, by selecting a role in the summary grid, and clicking the **Edit** icon in the Roles toolbar.

**Table 261: Access Control Roles Summary Grid**

| Option | Description |
|---|---|
| Roles toolbar | To manage roles, use the toolbar icons.<br>• Click the **Add** icon. to add a user role, provide the name and description for the role, and assign permissions in the **Role Information** page.<br>**NOTE**<br>To assign all the available permissions to the selected role, select the **Select all Permissions** check box.<br>• Click the **Vertical Ellipses** to perform any one of the following actions:<br>– **Edit**. Edit the selected user role, and modify the details for the role in the **Edit Role** dialog box.<br>– **Delete**. Delete a user role.<br>– **Clone**. Clone the selected user role<br>– **Export**. Export a user role.<br>• Click the **Horizontal Ellipses** to perform any one of the following actions:<br>– **Delete**. Click delete to Delete a role.<br>– **Export**. Click export to export a role.<br>– Click **Import** to import a role from the **Import Roles** page. You can import existing roles, overwrite existing roles, and skip roles. |
| Filters | Activates you to search the list of roles according to the following criteria:<br>• Role Name<br>• Role Description<br>• Modified By |

You can view details for the user accounts and user groups associated with a selected role in the Details panes.

**Table 262: Access Control Roles Details Panes**

| Option | Description |
|---|---|
| Role Name | Name of the role to apply to a specific level of users, such as user for base users or administrator for users with administrative permissions. |
| Description | Description of the role, indicating its purpose. |
| Modified By | Displays the name of the user who last updated the role. |
| Last Modified | Displays the date and time when the role was last updated. |
| Permissions | Displays the permissions assigned to the role according to the categories: Subscriptions, Administration, Automation Central, Configure, Data Sources, |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Environment, Home, Optimize, Plan, Troubleshoot, and Visualize. Expand the tree of each category to view all the assigned permissions.<br><br>You can edit the permissions assigned to the role by clicking the **Edit Role** icon.<br><br>• Click the **Expand All** button to expand the trees of all the categories, and select the check boxes to apply permissions for the selected role. |
| User Accounts | The roles are assigned to the users. You can only view users who have been assigned roles directly, the users belonging to a user group are not displayed here. The information in this pane is based on the data entered when you created the user, or imported with the user.<br>• First Name. Indicates the first name of each user who is assigned this role.<br>• Last Name. Indicates the last name of each user who is assigned this role.<br>• User name , without spaces, that will log in to VMware Cloud Foundation OperationsVMware Aria Operations<br>• Email. Indicates the email address for each user who is assigned this role. |
| User Groups | The roles are assigned to the selected user groups.<br>• Group Name: Name of each group that is associated with the selected role.<br>• Members: Number of members in each group. |

### Add or Edit Roles and Assign Permissions

You can add or edit user-specific roles to perform actions and view features and objects in VMware Aria OperationsVMware Cloud Foundation Operations. With role-based access, you can assign roles to users and let them only perform the actions that their permissions allow.

Where You Add or Edit User Roles

1. To manage user roles, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Select the **Roles** tab and click **Add**.
3. To edit a role, select the role and then click **Edit Role**. You can also click the **EDIT ROLE** button in the **Role Details** page and edit the role.

**Table 263: Add or Edit Role Information Page**

| Option | Description |
|---|---|
| Name | Name of the role to apply to a specific level of users, such as user for base users or administrator for users with administrative permissions. |
| Description | Description of the role, indicating its purpose. |
| **Assign Permissions**<br><br>**NOTE**<br>You can set default permissions for the selected role using the **Apply Default Permissions** option. The Apply Default Option is available only when you edit an Out-of-the-box role. | |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Select all Permissions | Click **Select all Permissions** to apply the permissions for the role you create or select. You can click the arrow next to the categories to set specific permissions to the roles you define. |
| Expand All | Click the **Expand All** button to expand the trees of all the categories and select the check boxes to apply or modify the permissions for the selected role. |
| Collapse All | Click the **Collapse All** button to Close the tree view and display only the high-level categories. |

4. Click **Save**.

### Export and Import of Roles

You can export User Roles configurations from one VMware Aria OperationsVMware Cloud Foundation Operations and import into any VMware Aria OperationsVMware Cloud Foundation Operations.

Export User Roles

1. From the left menu click **Administration › Control Panel**, and then click the **Access Control** tile.
2. In the **Roles** tab, select the user roles to be exported. Click the horizontal ellipsis next to **Add** and then, click **Export**.

The roles `.json` file is downloaded to the default download location.

Import Roles

Invalid privilege keys of a role are ignored during import.

1. From the left menu click **Administration › Control Panel**, and then click **Access Control** tile.
2. In the **Roles** tab, click the horizontal ellipsis next to **Add** and then, click **Import**.
3. Click **Browse** and select the roles `.json` file.
4. In case of a conflict, select either **Overwrite existing Roles** or **Skip Roles**.
5. Click **Import**.

Important Points

- Invalid permissions are ignored during import based on target VMware Aria OperationsVMware Cloud Foundation Operations.
- An error message is displayed with details for failed imports.

### Access Control: Scopes Tab

You can add, edit, or remove scope associated with users and groups in VMware Aria OperationsVMware Cloud Foundation Operations. Scope provides you the ability to limit access of a user or a set of users to VMware Aria OperationsVMware Cloud Foundation Operations. The All Objects scope is created by default. Assign the All Objects scope to the user to access all objects discovered by VMware Aria OperationsVMware Cloud Foundation Operations.

#### Where to Manage Scope

1. To manage scope, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Click the **Scopes** tab.

To view the details for a specific scope defined in VMware Aria OperationsVMware Cloud Foundation Operations, click the scope. The scope details are displayed in the right-side panel. The scope details include the user accounts and user groups associated with the scope.

**Table 264: Access Control Scope Summary Grid**

| Option | Description |
|---|---|
| Scope Toolbar | To manage scopes, use the toolbar options.<br>• Click **Add** to add the scope and provide the name and description for the scope, and assign members to objects in the **Scope Information** page.<br>Select **Object Hierarchies** and then select the object to assign members to. You can select one or more objects.<br>**NOTE**<br>For vSphere Storage object hierarchy, select the **Propagate to children** checkbox to include all future discovered objects in the scope.<br>• Click the **Vertical Ellipses** to perform any one of the following actions:<br>– **Edit**. Edit a specific scope and modify the details for the scope in the **Edit Scope Information** page.<br>– **Clone**. Clone a scope and enter a name and description for the cloned scope.<br>– **Delete**. Delete a scope.<br>• Click the **Horizontal Ellipses** and select **Delete**. Delete a scope. |
| Filters | Activates you to search the list of scopes according to the following criteria:<br>• Scope Name<br>• Description<br>• Modified By |

**Table 265: Access Control Scope Details Grid**

| Option | Description |
|---|---|
| Scope Name | Displays the name of the selected scope. |
| Description | Displays the description of the scope, indicating its purpose. |
| Modified By | Indicates the last person to update the scope. |
| Last Modified | Indicates the time when the scope was last updated. |
| User Accounts | Displays the user accounts associated with the selected scope. |
| User Groups | Displays the user groups associated with the selected scope. |

### Modify Scopes and Assign Member Objects

You can add and modify the scope defined in VMware Aria OperationsVMware Cloud Foundation Operations. Scope lets you create a package and associate the objects and object hierarchies to the package. You can select the object hierarchy and assign the required objects to the selected object hierarchy. After you define a scope, you can associate the scope to user accounts and user groups.

Where You Add or Edit Scope

1. To manage scope, from the left menu, click **Administration › Control Panel**, and then click the **Access Control** tile.
2. Click the **Scopes** tab.

3. In the **Scopes** tab, click **Add**.
4. To edit a scope, select the scope and then click **Edit Scope**. You can also click the **EDIT SCOPE** button in the **Scope Details** grid and edit the scope.

   **Table 266: Access Control Scopes Details Grid**

| Option | Description |
|---|---|
| Scope Name | Name of the scope to apply to a specific user or a set of users, such as user for base users or administrator for users with administrative permissions. |
| Description | Description of the scope, indicating its purpose. |

5. To add or edit the scope, select the required **Object Hierarchy**, and then select the **Object** to be associated.
6. Click **Save**.

After you save the scope, you can assign the scope to user accounts and user groups.

### Authentication Sources

Authentication Sources
Authentication Sources
VMware Aria OperationsVMware Cloud Foundation Operations uses authentication sources that activate you to import and authenticate users and user group information that reside on another machine: the Lightweight Directory Access Protocol (LDAP) platform-independent protocol, Active Directory, VMware Identity Manager, VMware SSO, Single Sign-On, and Others.

#### Where You Manage Authentication Sources

To manage authentication sources, from the left menu, click **Administration › Control Panel**, and then click the **Authentication Sources** tile.

**Table 267: Authentication Sources Toolbar and Data Grid**

| Option | Description |
|---|---|
| Authentication Sources toolbar | To manage authentication sources, use the toolbar icons.<br>• **Add** icon: Add an authentication source, and provide the information for the source in the Add Source for User and Group Import dialog box.<br>• Click the **Vertical Ellipses** to perform any one of the following actions:<br>– **Edit**. Edit the selected authentication source, and modify the details in the Edit Source dialog box.<br>– **Delete**. Delete an authentication source.<br>– **Synchronize User Groups**. Synchronize users within the groups imported through the selected Active Directory or LDAP authentication source. |
| Source Display Name | Name that you assign to the authentication source. |
| Source Type | Indicates the type of directory services access technology to access the source machine where the authentication database of user accounts resides. Options include:<br>• SSO SAML: An open-standard data format that activates Web browser single sign-on.<br>• VMware Identity Manager: A platform where you can manage users and groups, manage resources and user authentication, and access policies and entitle users to resources. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • VMware SSO: Browser based SSO that uses an external identity provider to log into vCenter Server, VMware Aria Operations, VMware Aria Operations for Logs, and VMware Aria Operations Orchestrator without providing credentials again.<br>• Open LDAP: A platform-independent protocol that provides access to an LDAP database on another machine to import user accounts.<br>• Active Directory or Other: Specifies any other LDAP-based directory services, such as Novel or Open DJ, used to import user accounts from an LDAP database on a Linux Mac machine. |
| Host | Name or IP address of the host machine where the user database resides. |
| Port | Port used for the import. |
| Base DN | Base distinguished name for the user search. VMware Aria OperationsVMware Cloud Foundation Operations locates only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although VMware Aria OperationsVMware Cloud Foundation Operations populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration. |
| Auto Synchronization | When selected, activates VMware Aria OperationsVMware Cloud Foundation Operations to map imported LDAP users to user groups. |
| Last Synchronized | Date and time that the synchronization last occurred. |

### Configure VMware Single Sign-On for VMware Aria Operations

After installing or upgrading to VMware Aria Operations 8.18, you can configure VMware Aria Operations for VMware Single Sign-On. When you configure VMware Single Sign-On, you use an external identity provider to sign into VMware Aria Operations.

- Ensure that the associated vCenter Server host is configured for VMware Single Sign-On. For more information about configuring a vCenter Server host for VMware Single Sign-On, see Configure VMware Single Sign-on.

  **NOTE**
  After configuring VMware Single Sign-On configuration for VMware Aria Operations, you can still log in to VMware Aria Operations with a local account.

1. Log in to VMware Aria Operations with a local account. For more information, see Logging In to VMware Aria Operations.
2. Configure VMware SSO as an authentication type. For more information, see Authentication Sources and Authentication Sources: Add Authentication Source for User and Group Import.

   **NOTE**
   Only one VMware Single Sign-On configuration can exist during any period of time.

3. To allow users or groups to log in to VMware Aria Operations using **VMware SSO** authentication source, import users or groups from the authentication source into VMware Aria Operations. For more information, see Import User Accounts From Source and Import User Groups From Source.

   **NOTE**
   You must select a role for each user and assign a scope for each role.

4. A certificate of type **VMware SSO** is created and can be viewed from **Administration › Control Panel › Certificates**.
5. Users from the VMware SSO authentication source can now log into VMware Aria Operations using VMware SSO. When you log in to VMware Aria Operations using VMware SSO, you will be redirected to an external authentication page. Enter the credentials to log in to VMware Aria Operations.

**NOTE**
- You can delete and unregister a VMware Single Sign-On authentication source. For more information see, Delete and Deregister a VMware Single Sign-On Authentication Source.
- You have to re-register the VMware SSO server in certain scenarios. Navigate to the **Authentication Sources** page, click the **Re-register source** link and enter the vCenter Server credentials. Re-register the VMware SSO server in the following two scenarios:
  – If you add or remove nodes from a cluster in VMware Aria Operations, and/or
  – If you have modified the system access URL when a load balancer is used (**Administration › Global Settings › System Settings › System Access URL option**).

### Authentication Sources: Add Authentication Source for User and Group Import

When you import user account information that resides on another machine, you must define the criteria used to import the user accounts from the source machine.

#### Where You Add or Edit Authentication Sources

1. To add authentication sources, from the left menu, click **Administration › Control Panel**, and then click the **Authentication Sources** tile.
2. Click **Add**.
3. To edit authentication sources, click **Edit**.

**Table 268: Authentication Sources Add Source for User and Group Import**

| Option | Description |
|---|---|
| Source Display Name | Name that you assign to the authentication source. |
| Source Type<br><br>**NOTE**<br>The option you select in the **Source Type** drop-down box, determines the options available in this dialog box. | Indicates the type of directory services access technology to access the source machine where the database of user accounts resides. There are two types of databases: LDAP and single sign-on. Options include:<br>• SSO SAML: An XML-based standard for a web browser single sign-on that activates users to perform single sign-on to multiple applications.<br>• VMware Identity Manager: A platform where you can manage users and groups, manage resources and user authentication, and access policies and entitle users to resources.<br>• VMware SSO: Browser based SSO that uses an external identity provider to log into vCenter Server, VMware Aria Operations, VMware Aria Operations for Logs, and VMware Aria Operations Orchestrator without providing credentials again. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Open LDAP: A platform-independent protocol that provides access to an LDAP database on another machine to import user accounts.<br>• Other: Specifies any other LDAP-based directory services, such as Novel or OpenDJ, used to import user accounts from an LDAP database on a Linux Mac machine. |

**Table 269: Authentication Sources Add Source for User and Group Import - Options Available When SSO SAML Is Selected.**

| Name | Description |
|---|---|
| Host | Name or IP address of the host machine where the single sign-on user server resides. |
| Port | The single sign-on listening port. By default this is set to 443. |
| User Name | Name of the user account that can log in to the single sign-on host machine. |
| Password | Password of the user account that can log in to the single sign-on host machine. |
| Grant administrator role to VMware Aria OperationsVMware Cloud Foundation Operations for future configuration? | When you create a single sign-on source, a new VMware Aria OperationsVMware Cloud Foundation Operations user account is created on the single sign-on server.<br>• Select **Yes**, to grant VMware Aria OperationsVMware Cloud Foundation Operations an administrative role so that it can be used to configure the SSO source if changes are made to the VMware Aria OperationsVMware Cloud Foundation Operations setup.<br>• If you select **No** and the VMware Aria OperationsVMware Cloud Foundation Operations setup is changed, SSO users will not be able to log in until you re-register the SSO source. |
| Automatically redirect to VMware Aria OperationsVMware Cloud Foundation Operations single sign-on URL? | After you have configured a single sign-on source, users are redirected to the vCenter SSO server.<br>• Select **Yes**, to redirect users to the single sign-on server for authentication.<br>• If you select **No** users must sign in through the VMware Aria OperationsVMware Cloud Foundation Operations login page. |
| Import single sign-on user groups after adding the current source? | When you have set up a single sign-on source, you import users and user groups into VMware Aria OperationsVMware Cloud Foundation Operations so that single sign-on users can access the system with their single sign-on permissions.<br>• If you select **Yes**, the wizard directs you to the Import User Groups page so that you can import user groups when you have finished setting up the SSO source. |

*Table continued on next page*

*Continued from previous page*

| Name | Description |
|---|---|
| | • If you want to import user accounts, or user groups at a later stage, select **No**. |
| Advanced | If your system uses a load balancer, enter the IP address of the load balancer. |
| Test | Tests whether the host machine can be reached with the credentials provided. |

**Table 270: Authentication Sources Add Source for User and Group Import - Options Available When VMware SSO Is Selected.**

| Name | Description |
|---|---|
| Host | Name or IP address of the vCenter Server host machine on which VMware SSO has been configured. |
| Port | The single sign-on listening port. By default this is set to 443. |
| Tenant | |
| Username | Name of the vCenter Server user account that can log in to the VMware SSO host machine.<br>**NOTE**<br>The user must have the *`VcIdentityProviders.Manage`* permission assigned. |
| Password | Password of the vCenter Server user account that can log in to the VMware SSO host machine. |

**NOTE**
Only one VMware Single Sign-On configuration can exist during any period of time.

**Table 271: Authentication Sources Add Source for User and Group Import - Options Available When Open LDAP, Active Directory, and Other Are Selected.**

| Option | Description |
|---|---|
| Integration Mode Basic settings | Applies basic settings to integrate the LDAP import source with the instance of VMware Aria OperationsVMware Cloud Foundation Operations.<br><br>Use Basic integration mode to have VMware Aria OperationsVMware Cloud Foundation Operations discover the host machine where the LDAP database resides, and set the base distinguished name (Base DN) used to search for users. You provide the name of the domain and the subdomain, which VMware Aria OperationsVMware Cloud Foundation Operations uses to populate the Host and Base DN details, and the name and password of the user who can log in to the LDAP host machine.<br><br>In Basic mode, attempts to fetch the host and port from the DNS server, and obtain the Global Catalog and domain controllers for the domain, with preference given to SSL/TLS-activated servers. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Domain/Subdomain. Domain information for the LDAP user account.<br><br>**NOTE**<br>To import users and groups from multiple subdomains, use the root domain.com instead of subdomain. Using a subdomain limits the visibility of VMware Aria OperationsVMware Cloud Foundation Operations to groups and users from that specific subdomain.<br><br>• Use SSL/TLS. When selected, VMware Aria OperationsVMware Cloud Foundation Operations uses the Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocol to provide secure communication when you import users from an LDAP database. You do not need to install the SSL/TLS certificate. Instead, VMware Aria OperationsVMware Cloud Foundation Operations prompts you to view and verify the thumbprint, and accept the LDAP server certificate. After you accept the certificate, the LDAP communication proceeds.<br>• If Active Directory uses a self-signed certificate, then the certificate should contain the **Subject Alternative Name** field. VMware Aria OperationsVMware Cloud Foundation Operations can successfully verify the Active Directory certificate and integrate with Active Directory only if, the host name or the IP address provided in the **Subject Alternative Name** field matches the address of the domain controller on which the certificate is used.<br>• User Name. Name of the user account that can log in to the LDAP host machine.<br>• Reset Password. Reset the password of the user account that can log in to the LDAP host machine.<br>• Automatically synchronize user membership for configured groups. When selected, activates VMware Cloud Foundation OperationsVMware Aria Operations to map imported LDAP users to user groups.<br>• Host. Name or IP address of the host machine where the LDAP user database resides.<br>• Port. Port used for the import. Use port 389 if you are not using SSL/TLS, or port 636 if you are using SSL/TLS, or another port number of your choice. Global Catalog ports are 3268 for non-SSL/TLS, and 3269 for SSL/TLS.<br>• Base DN. Base distinguished name for the user search. VMware Aria OperationsVMware Cloud Foundation Operations locates only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although VMware Aria OperationsVMware Cloud Foundation Operations populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration.<br>• Common Name. LDAP attribute used to identify the user name. The default attribute for Active Directory is `userPrincipalName`. |
| Integration Mode Advanced settings | Applies advanced settings to integrate the LDAP import source with the instance of VMware Aria OperationsVMware Cloud Foundation Operations.<br><br>Use Advanced integration mode to manually provide the host name and base distinguished name (Base DN) to have VMware Aria OperationsVMware Cloud Foundation Operations import users. You provide the name and password of the user who can log in to the LDAP host machine. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Host. Name or IP address of the host machine where the LDAP user database resides.<br>• Use SSL/TLS. When selected, VMware Aria OperationsVMware Cloud Foundation Operations uses the Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocol to provide secure communication when you import users from an LDAP database. You do not need to install the SSL/TLS certificate. Instead, VMware Aria OperationsVMware Cloud Foundation Operations prompts you to view and verify the thumbprint, and accept the LDAP server certificate. After you accept the certificate, the LDAP communication proceeds.<br>• If Active Directory uses a self-signed certificate, then the certificate should contain the **Subject Alternative Name** field. VMware Aria OperationsVMware Cloud Foundation Operations can successfully verify the Active Directory certificate and integrate with Active Directory only if, the host name or the IP address provided in the **Subject Alternative Name** field matches the address of the domain controller on which the certificate is used.<br>• Base DN. Base distinguished name for the user search. VMware Cloud Foundation OperationsVMware Aria Operations will locate only the users under the Base DN. The Base DN is an elementary entry for an imported user's distinguished name (DN), which is the base entry for the user name without the need for other related information such as the full path to the user account, or the inclusion of related domain components. Although VMware Cloud Foundation OperationsVMware Aria Operations populates the Base DN, an Administrator must verify the Base DN before saving the LDAP configuration.<br>• User Name. Name of the user account that can log in to the LDAP host machine.<br>• Reset Password. Reset the password of the user account that can log in to the LDAP host machine.<br>• Automatically synchronize user membership for configured groups. When selected, activates VMware Aria OperationsVMware Cloud Foundation Operations to map imported LDAP users to user groups.<br>• Common Name. LDAP attribute used to identify the user name. The default attribute for Active Directory is `userPrincipalName`.<br>• Port. Port used for the import. Use port 389 if you are not using SSL/TLS, or port 636 if you are using SSL/TLS, or another port number of your choice. Global Catalog ports are 3268 for non-SSL/TLS, and 3269 for SSL/TLS. |
| Search Criteria | Displays the search criteria settings.<br><br>Although VMware Aria OperationsVMware Cloud Foundation Operations populates part of the search criteria, an Administrator must verify the settings to ensure that the settings are correct according to the properties of the LDAP type.<br><br>• Group Search Criteria. Search criteria to find LDAP groups. If not included, VMware Cloud Foundation OperationsVMware Aria Operations uses the default search parameters: `(|(objectClass=group)(objectClass=groupOfNames))`<br>• Member Attribute. Name of the attribute for a group object that contains the list of members. If not included, VMware Cloud Foundation OperationsVMware Aria Operations uses member by default.<br>• User Search Criteria. Search criteria to use the member field to find and cache LDAP users. You enter sets of key=value pairs in the form `(|(key1=value1)(key2=value2))`. If not included, VMware Cloud Foundation OperationsVMware |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Aria Operations searches for each user separately. This operation might take extra time.<br>• Member Match Field. Name of the attribute for a user object to match with the member entry from a group object. If not included, VMware Cloud Foundation OperationsVMware Aria Operations treats the member entry as a distinguished name.<br>• LDAP Context Attributes. Attributes that VMware Cloud Foundation OperationsVMware Aria Operations applies to the LDAP context environment. You enter sets of key=value pairs separated by commas, such as **`java.naming.referral=ignore,java.naming.ldap.deleteRDNfalse`**. |
| Test | Tests whether the host machine can be reached, with the credentials provided. Although a test of the connection is successful, users who use the search feature must have read permissions in the LDAP source.<br><br>This test does not verify the accuracy of the Base DN or Common Name entries. |

**Table 272: Authentication Sources Add Source for User and Group Import - Options Available When VMware Identity Manager Is Selected.**

| Option | Description |
|---|---|
| Host | Name or IP address of the VMware Identity Manager machine where the single sign-on user server resides. |
| Port | The single sign-on listening port. By default this is set to 443. |
| Tenant | This is an optional field. |
| User name | VMware Identity Manager system-domain tenant administrator user name. |
| Password | Password of the VMware Identity Manager system-domain tenant administrator. |
| Redirect IP/ FQDN | This is the IP address of VMware Cloud Foundation OperationsVMware Aria Operations node where a user is redirected after a successful authentication from VMware Identity Manager. By default, this is the IP address of the VMware Cloud Foundation OperationsVMware Aria Operations primary node.<br><br>**NOTE**<br>When the primary replica becomes the primary node on VMware Cloud Foundation OperationsVMware Aria Operations, then VMware Cloud Foundation OperationsVMware Aria Operations administrator has to manually edit the IP address and set it to the IP address of the current primary node. |
| Test | Tests whether the VMware Identity Manager machine can be reached, with the credentials provided. |

### Export and Import of Authentication Sources

You can export Authentication source configurations from one VMware Aria OperationsVMware Cloud Foundation Operations and import into any VMware Aria OperationsVMware Cloud Foundation Operations.

#### Export Authentication Sources

1. From the left menu click **Administration › Control Panel**, and then click **Authentication Sources** tile.
2. In the **Authentication Sources** tab, select the authentication sources to be exported. Click the horizontal ellipsis next to **Add** and then, click **Export**.
3. You are prompted to enter a password when you export authentication sources. Enter the password and note it down, you have to use the same password when you import the authentication sources.
4. Click **Export**.

The authentication sources `.json` file is downloaded to the default download location.

#### Import Authentication Sources

1. From the left menu click **Administration › Control Panel**, and then click **Authentication Sources** tile.
2. In the **Authentication Sources** tab, click the horizontal ellipsis next to **Add** and then, click **Import**.
3. Click **Browse** and select the authentication sources `.json` file.
4. Enter the same password which you had used during authentication sources export.
5. In case of a conflict, select either **Overwrite existing Authentication Sources** or **Skip Authentication Sources**.
6. Click **Import**.
7. After the import you must re-accept all untrusted certificates manually.

#### Important Points

- An error message is displayed with details for failed imports.
- Export or Import of authentication sources is not supported in VMware Aria Operations (SaaS).

### Delete and Deregister a VMware Single Sign-On Authentication Source

Deleting the VMware Single Sign-on authentication source automatically removes users and user groups that are associated with the VMware Single Sign-On authentication source.

- A VMware SSO authentication source must be configured.

1. From the left menu, click **Administration › Control Panel**, and then click the **Authentication Sources** tile.
2. Select the **VMware SSO** authentication source, and then click the vertical ellipsis and select **Delete**.
3. Enter the **User Name** and **Password** of the vCenter Server user account that can log in to the VMware SSO host machine.
4. Select **Force Delete** if you cannot access the single sign-on server. VMware Aria Operations continues to attempt to clean the VMware Single Sign-On registration even when errors occur.
5. Click **OK**.

### Audit Users and the Environment in VMware Aria OperationsVMware Cloud Foundation Operations

Audit Users and the Environment
Audit Users and the Environment

At times, you might need to provide documentation as an evidence of the sequence of activities that took place in your VMware Aria OperationsVMware Cloud Foundation Operations environment. Auditing allows you to view the users, objects, and information that is collected. To meet audit requirements, such as for business critical applications that contain sensitive data that must be protected, you can generate reports on the activities of your users, the privileges assigned to users to access objects, and the counts of objects and applications in your environment.

Auditing reports provide traceability of the objects and users in your environment.

**User Activity Audit**

Run this report to understand the scope of user activities, such as logging in, actions on clusters and nodes, changes to system passwords, activating certificates, and logging out.

**User Permissions Audit**

Generate this report to understand the scope of user accounts and their roles, access groups, and access privileges.

**System Audit**

Run this report to understand the scale of your environment. This report displays the counts of configured and collecting objects, the types and counts of adapters, configured and collecting metrics, super metrics, applications, and existing virtual environment objects. This report can help you determine whether the number of objects in your environment exceeds a supported limit.

**VC Pricing Migration Audit**

Run this report to understand the details of VC pricing migration.

**Chargeback Migration Audit**

Run this report to understand the details of Chargeback migration.

**System Component Audit**

Run this report to display a version list of all the components in your environment.

## Reasons for Auditing Your Environment

Auditing in VMware Aria OperationsVMware Cloud Foundation Operations helps data center administrators in the following types of situations.

- You must track each configuration change to an authenticated user who initiated the change or scheduled the job that performed the change. For example, after an adapter changes an object, which is associated with a specific object identifier at a specific time, the data center administrator can determine the principal identifier of the authenticated user who initiated the change.
- You must track who made changes to your data center during a specific range of time, to determine who changed what on a particular day. You can identify the principal identifiers of authenticated users who were logged in to VMware Aria OperationsVMware Cloud Foundation Operations and running jobs, and determine who initiated the change.
- You must determine which objects were affected by a particular user during a time-specific range of time.
- You must correlate events that occurred in your data center, and view these events overlayed so that you can visualize relationships and the cause of the events. Events can include login attempts, system start up and shutdown, application failures, watchdog restarts, configuration changes of applications, changes to security policy, requests, responses, and status of success.
- You must validate that the components installed in your environment are running the latest version.

### User Activity Audit

The user activity report helps you understand the scope of user activities in your VMware Aria OperationsVMware Cloud Foundation Operations instance, such as when users logged in, actions they took on clusters and nodes, changes they made to system passwords, when they activated certificates, and when they logged out.

### Where You Audit User Activity

To audit user activity, from the left menu, click **Administration › Control Panel**, and then click the **Audit** tile. The activities that users performed in the environment appear on the page.

**Table 273: User Activity Audit Actions**

| Option | Description |
|---|---|
| Download | Download the user activity audit information to a report in PDF or XLS format. |
| Configure | Configure the settings to send the user activity log to an external syslog server to meet security auditing requirements.<br>• Output log to external syslog server. When selected, VMware Aria OperationsVMware Cloud Foundation Operations sends the log to a separate server machine.<br>• IP Address or Host Name. Identification for the syslog server.<br>• Port. VMware Aria OperationsVMware Cloud Foundation Operations port used to send the audit information to the external server. |
| Date Range | Display the list of user activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times. |
| Starting Line | Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |
| Number of Lines | Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0. |
| Filter | Filters the data according to User ID, User Name, Auth Source, Session, Client IP, Category, and Message. |

### User Permissions Audit

A user permissions audit report provides an overview of the local users and LDAP imported users in your VMware Aria Operations instance, and a list of groups to which each user belongs. This report helps you understand the scope of the user accounts and their roles, access groups, and access privileges in your environment.

A user permissions audit report provides an overview of the local users in your VMware Aria OperationsVMware Cloud Foundation Operations instance, and a list of groups to which each user belongs. This report helps you understand the scope of the user accounts and their roles, access groups, and access privileges in your environment.
The report displays the access group associated with each local user and LDAP imported user and the access privileges granted to the user in each access group. This report does not include vCenter users, roles, or privileges.

The report displays the access group associated with each local user and the access privileges granted to the user in each access group. This report does not include vCenter users, roles, or privileges.

When a user is a member of a specific user group, the associated access group could provide the user with access to configuration, dashboards, and templates, or to specific navigation areas in the user interface such as Administration. The access rights associated with the access group include actions for each access group, such as the ability to add, edit, or delete dashboards, or to view, configure, or manage objects.

### Where You Audit User Permissions

1. To audit user permissions, from the left menu, click **Administration › Control Panel**, and then click the **Audit** tile.
2. Click the **User Permissions Audit** tab.

The permissions assigned to users, and their associated access groups and access privileges, appear on the page.

**Table 274: User Permissions Audit Actions**

| Option | Description |
|---|---|
| Download | Download the user permissions audit information to a report in PDF or XLS format. |

### System Audit for VMware Aria OperationsVMware Cloud Foundation Operations

A system audit report provides an overview of the counts of objects, metrics, super metrics, applications, and custom groups in your VMware Aria OperationsVMware Cloud Foundation Operations instance. This report can help you understand the scale of your environment.

The system audit report displays the types and number of objects that VMware Aria OperationsVMware Cloud Foundation Operations manages. Reported objects include those that are configured and collecting data, the types of objects, object counts for adapters, the metrics that are configured and being collected, super metrics, VMware Aria OperationsVMware Cloud Foundation Operations generated metrics, the number of applications used, and the number of custom groups.

You can use this report to help determine whether the number of objects in your environment exceeds a supported limit.

#### Where You Audit the System

1. To audit the objects, metrics, applications, and custom groups in your environment, from the left menu, click **Administration › Control Panel**, and then click the **Audit** tile.
2. Click the **System Audit** tab.

The objects and their associated counts appear in the report.

**Table 275: System Audit Actions**

| Option | Description |
|---|---|
| Download | Download the system information to a report in PDF or XLS format. |

### System Component Audit

A system component audit report provides a version list of every component installed in the system.

#### Where You Audit System Components

1. To audit system components, from the left menu, click **Administration › Control Panel**, and then click the **Audit** tile.
2. Click the **System Component Audit** tab.

A list of components installed in the environment appears on the page.

**Table 276: System Component Audit Actions**

| Option | Description |
|---|---|
| Download | Display the version information in a new browser window. |

### VC Pricing Migration Audit

The vCenter Pricing Migration Audit provides details of the migration of pricing cards to VMware Aria OperationsVMware Cloud Foundation Operations policies.

#### Where You Audit vCenter Pricing Migration Components

1. To audit vCenter Pricing Migration components, from the left menu, click **Administration › Control Panel** and then click the **Audit** tile.
2. From the **Audit** page, click **Chargeback Audit tab › VC Pricing Migration Audit tab**.

**Table 277: vCenter Pricing Migration Audit Actions**

| Option | Description |
|---|---|
| Download | Download the vCenter Pricing migration information to a report in PDF or XLS format. |
| Date Range | Display the list of vCenter Pricing migration activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times. |
| Starting Line | Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |
| Number of Lines | Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0. |
| Filter | Filters the data according to User ID, User Name, Auth Source, Session, Client IP, Category, and Message. |

**NOTE**
For more information about vCenter pricing migration, see Migration of vCenter Pricing Cards to vCenter Pricing Policies

### Chargeback Migration Audit

The Chargeback Migration Audit provides details of the migration of VMware Chargeback configurations and content to VMware Aria OperationsVMware Cloud Foundation Operations.

#### Where You Audit Chargeback Migration Components

1. To audit Chargeback Migration components, from the left menu, click **Administration › Control Panel**, and then click the **Audit** tile.

2. From the **Audit** page, click **Chargeback Audit tab › Chargeback Migration Audit tab**.

**Table 278: Chargeback Migration Audit Actions**

| Option | Description |
|---|---|
| Download | Download the Chargeback migration information to a report in PDF or XLS format. |
| Date Range | Display the list of Chargeback migration activities performed in the past based on a selected number of hours, days, weeks, months, or years, or between two specific dates and times. |
| Starting Line | Indicates the starting line of the file . 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |
| Number of Lines | Specifies the number of lines to be displayed in the search result. For example: If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0. |
| Filter | Filters the data according to User ID, User Name, Auth Source, Session, Client IP, Category, and Message. |

**NOTE**
For more information about Chargeback migration, see Migration of Chargeback Data to VMware Aria OperationsVMware Cloud Foundation Operations in Case of an Upgrade .

### User Preferences in VMware Aria OperationsVMware Cloud Foundation Operations

User Preferences
User Preferences
You can configure the user preferences to determine the VMware Aria OperationsVMware Cloud Foundation Operations display options, such as the number of metrics and groups to display and whether to synchronize system time with the host machine.

To configure the user preferences, in the menu, click the  icon, and then click **Preferences**. The user preference settings appear in the dialog box.

**Table 279: User Preference Settings**

| Option | Description |
|---|---|
| Display | Configure how many metrics and root cause groups to display.<br>• Color scheme: Set the user interface to display in light or dark colors. |
| Language and Time | • Language: Select the language you want VMware Aria OperationsVMware Cloud Foundation Operations to be displayed in. You can select the browser language or select a language from the drop-down list.<br>• Font: Select the font for reports. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Synchronize the time used for the VMware Aria OperationsVMware Cloud Foundation Operations instance, and display the updated time when VMware Aria OperationsVMware Cloud Foundation Operations communicates with the host machine.<br>• Browser time. All dates and times displayed in the user interface use the time zone settings of the local browser.<br>• Host time. All dates and times displayed in the user interface use the time zone of the host machine.<br>• Show update time in the application header. Displays the updated time in the top-level header of the VMware Aria OperationsVMware Cloud Foundation Operations user interface. The updated timestamp appears to the left of the refresh button. Other features, such as dashboards, use the updated time to display data at specific intervals. |
| Product Survey Opt Out | Allows you to opt out of providing regular feedback from the VMware Aria Operations UI. |
| Account | Change the password for the user account. |

## VMware Aria OperationsVMware Cloud Foundation Operations Certificates

Certificates
VMware Aria OperationsVMware Cloud Foundation Operations includes a central page where you can review authentication certificate contents.

### How the Certificates Page Works

The Certificates page lets you examine certificate contents without the need to open the certificate outside of VMware Aria OperationsVMware Cloud Foundation Operations.

### Where You Find Certificates

In the menu, click **Administration › Control Panel**, and then click the **Certificates** tile.

### Certificate Tabs

The certificate tab describes columns of exceptions tabs.

> **NOTE**
> The CRL tab is activated only when you select the **Activate Standard Certificate Validation** under **Global Settings**.

**Table 280: Certificate Tabs**

| Tabs | Description |
|---|---|
| Exceptions | Lists the certificate that is accepted by the VMware Aria OperationsVMware Cloud Foundation Operations administrator but is not certified by the Certificate Authority (CA). |
| CRL | A Certificate Revocation List (CRL) is a list of digital certificates that have been revoked by the issuing |

*Table continued on next page*

*Continued from previous page*

| Tabs | Description |
|---|---|
| | Certificate Authority (CA) before their scheduled expiration date and should no longer be trusted. Click the Add icon to upload the certificates. |

**Certificate Options**

The options include a data grid for examining certificate contents.

**Table 281: Certificate Options**

| Option | Description |
|---|---|
| Certificate Thumbprint | Unique alphanumeric string associated with the certificate |
| Issued By | Content associated with the issuer of the certificate, such as organization name and location |
| Issued To | Typically, content associated with the issuer, plus the certificate object Identifier (OID) |
| Expires | The date after which the certificate cannot be used for successful authentication |

### Importing CA Certificates

Certificate Authority (CA) or root certificates are used for establishing the outgoing connections from VMware Aria OperationsVMware Cloud Foundation Operations. CA Certificates imported by the users will be used in the following VMware Aria OperationsVMware Cloud Foundation Operations domains: Authentication Sources (Active Directory (AD), Open LDAP, VMware Identity Manager), Outbound Plugins, and Adapter Endpoint.

1. In the menu, click **Administration › Control Panel**, and then click the **Certificates** tile.
2. Click **Import**.
   The Import CA Certificate(s) dialog box appears. You can only import certificates that are encoded in the PEM format.
3. Click **Browse**.
4. Locate the certificate `.pem` file and click **Open** to load the file in the Import CA Certificate(s) dialog box.
   The certificate information box appears with the certificate thumbprint, issued by, issued to and expiry date. For example, if you select a certificate that will expire in 10 days, you will receive a notification that the certificate is expiring soon.

   **NOTE**
   If a certificate is close to its expiry date, a corresponding notification is displayed on the **Home** page.

5. Click **Import**.
6. Click the **Vertical Ellipsis** to delete a certificate.

### Removing an Adapter Certificate

You can delete an old or expired certificate associated with an adapter.

1. In a Web browser, navigate to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://node-FQDN-or-ip-address/ui.
2. Log in with the administrator user name and password.

3. In the menu, click **Administration › Control Panel**, and then click the **Certificates** tile.
4. In the certificate window, select the certificate that has to be removed.
5. Click **Delete** to remove the certificate.

   **NOTE**
   You must own the credential used to configure the adapter to be able to delete the certificate. For more information, see Configuring Credentials in Integrations.

6. If the certificate is being used by the adapter, then the following message comes up:

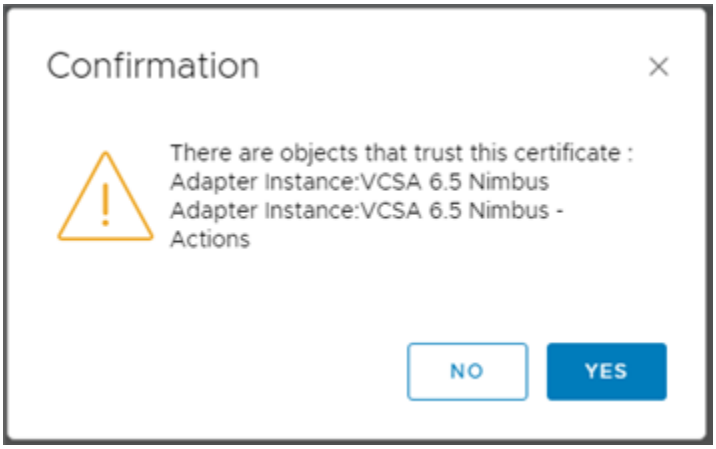


   A certificate can be configured for one or more adapters if it is the same destination system.

7. If you delete a certificate which is already being used by another adapter, the adapter fails to connect or start. As a workaround, perform the following steps:
   a) On the left pane, click **Administration › Integrations**.
      The Accounts tab appears.
   b) Select the particular integration and click the vertical ellipsis and then click **Edit**.
      The Account Information page opens.
   c) Click **Validate Connection**.

      **NOTE**
      If the adapter instance is saved using a credential that is owned by another user, you must apply your own credentials to validate the connection. For more information, see Manage Credentials.

   d) A prompt comes up asking the user to import the associated certificate. Click **OK**.
   e) Restart the adapter from the **Accounts** tab.

## VMware Aria OperationsVMware Cloud Foundation Operations Support Logs for Product UI

Logs
Use the expandable tree of VMware Aria OperationsVMware Cloud Foundation Operations log files for troubleshooting issues in the product UI.

### How VMware Aria OperationsVMware Cloud Foundation Operations Support Logs Work

You can browse and load the VMware Aria OperationsVMware Cloud Foundation Operations log files for review and for troubleshooting issues in the product UI.. You can also edit the log file folders, limit the retained log size, and set logging levels.

VMware Aria OperationsVMware Cloud Foundation Operations logs are categorized by cluster node, and log type. All logs are in the UTC formatted date and time. The logging format is as follows:

```
Date/Time+0000, LEVEL, [THREAD/IP Address], [Specific Fields], CLASS - MESSAGE
```

If you have configured a timezone for the VMware Aria Operations VM, the system logs will be in that timezone. The VMware Aria Operations logs will remain in UTC.

**Where You Find VMware Aria OperationsVMware Cloud Foundation Operations Support Logs**

In the menu, click **Administration › Control Panel**, and then click the **Support Logs** tile.

**Log Viewer Options**

Use the toolbar options to control the tree of items and the viewer.

1. Click **Node** and select any component that is listed under the node.
2. Click the gear icon, enter the logging levels and log size.
3. Click **OK**.

> **NOTE**
> Not all components have relevant syslog information. Therefore, not all nodes have the configuration option activated.

**Figure 19: Logs**

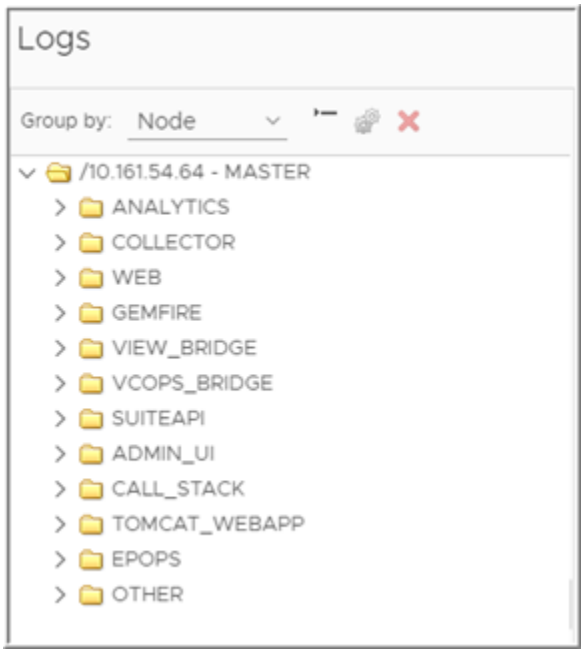


**Figure 20: Log Options**

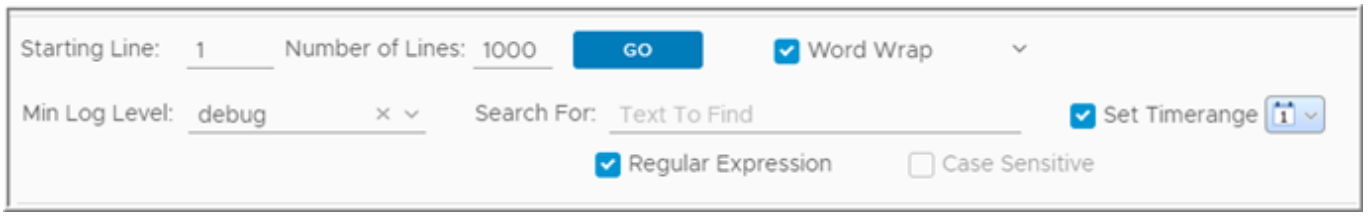

**Table 282: Log Viewer Toolbar Options**

| Option | Description |
|---|---|
| Group By | Organizes the tree by cluster node or log type. |
| Collapse All | Closes the view of the tree to show only the high-level folders. |
| Edit Properties | For the selected folder, you can limit the log size and set logging levels. |
| Delete Selected File | Deletes the log file. |
| Starting Line | Indicates the starting line of the file .<br>0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |
| Number of Lines | Specifies the number of lines to be displayed in the search result.<br>For example:<br><br>If you want to see the first 10 occurrences of a particular chunk of text, enter the number of lines as 10 and the starting line as 0. |
| Min Log Level | If you specify the minimum log level, the logs for that particular log level and higher are shown.<br>For example:<br><br>If you select **warning**, the logs having the same log level (**warning**) and higher are shown . |
| Text to Find | Enter the specific text that you want to search in the logs. Add the following filters for search, if required:<br>• **Case Sensitive**<br>• **Regular Expression**<br>You can perform the search at various levels:<br><br>• On a single file: Use this option if you want to search a single log file .<br>• On all the log files of an entity: Use this option if you want to search all the log files of an entity such as a log type or folder.<br>• On all the log files of a node: Use this option if you want to search all the log files that are grouped under a node.<br>The last modified time for any file is found by placing the pointer on the file in the tree. |
| Set Timerange | If you specify a time range, the logs for that particular time range are shown in the search results. |
| Word Wrap | If you select this option, the part of the line that does not fit on the screen is moved to the next line. If you do not select this option, a scroll bar is provided to see the complete line. |

## Cluster Management

VMware Aria Operations includes a central page where you can monitor and manage the nodes in your VMware Aria Operations cluster and the adapters that are installed on the nodes.

### How Cluster Management Works

Cluster management lets you view and change the online or offline state of the overall VMware Aria Operations cluster or the individual nodes. In addition, you can activate or deactivate high availability (HA) and view statistics related to the adapters that are installed on the nodes.

**Where You Find Cluster Management**

From the left menu, click **Administration › Control Panel**, and then click the **Cluster Management** tile.

**Cluster Management Options**

The options include cluster-level monitoring and management features.

**Table 283: Initial Setup Status Details**

| Option | Description |
|---|---|
| Actions | Expand Actions.<br>• Click **Network Time Protocol Settings** to set the global network time protocol settings for your cluster. For more information see the 'Configuring the Global Network Time Protocol (NTP) Settings' topic in the *VMware Aria Operations Configuration Guide*.<br>• Click **Rebalance** to rebalance the adapter, disk, memory, or network load across VMware Aria Operations cluster nodes to increase the efficiency of your environment. |
| Cluster Status | Displays the online, offline, or unknown state of the VMware Aria Operations cluster.<br>Once CA is activated, it displays the status of the two fault domains. |
| High Availability | Indicates whether HA is activated, deactivated, or degraded. |
| Continuous Availability | Indicates whether CA is activated, deactivated, or degraded. |

VMware Aria Operations provides node-level information and a toolbar for taking nodes online or offline.

**Table 284: Nodes in the VMware Aria Operations Cluster**

| Option | Description |
|---|---|
| Node Name | Machine name of the node.<br>The node that you are logged into displays a dot next to the name. |
| Node Address | Internet protocol (IP) address of the node. Primary and replica nodes require static IP addresses. Data nodes can use DHCP or static IP. |
| Cluster Role | Type of VMware Aria Operations node: primary, data, replica, or witness node (only applicable for CA). |
| Fault Domain | Displays the fault domain a node is associated to in a CA activated cluster.<br>**NOTE**<br>This column appears only if CA is activated. |
| Node Pair | Displays which pair the node belongs to. For example, in CA, nodes are added in pairs. If there are four nodes, the |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | column displays whether the node is part of pair one or two.<br>**NOTE**<br>This column appears only if CA is activated. |
| State | Running, Not Running, Going Online, Going Offline, Inaccessible, Failure, Error |
| Status | Online, offline, unknown, or other condition of the node. |
| Objects in Process | Total environment objects that the node currently monitors. |
| Objects Being Collected | Total environment objects that the node collected. |
| Metrics in Process | Total metrics that the node has discovered since being added to the cluster. |
| Metrics Being Collected | Total metrics the node has collected since being added to the cluster. |
| Version | Displays the VMware Aria Operations software version and the build number installed on the node. |

In addition, there are adapter statistics for the selected node.

**Table 285: Adapters on Server**

| Option | Description |
|---|---|
| Name | Name that the installing user gave to the adapter. |
| Status | Indication of whether the adapter is collecting data or not. |
| Objects Being Collected | Total environment objects that the adapter currently monitors. |
| Metrics Being Collected | Total metrics that the adapter has collected since being installed on the node. |
| Last Collection Time | Date and time of the most recent data collection by the adapter. |
| Added On | Date and time when the adapter was installed on the node. |

## Monitoring Data Collection

Monitoring Data Collection
The collection status page provides an overview of the data that is being collected at the cluster level in VMware Aria OperationsVMware Cloud Foundation Operations. You can view the details for each collector and you can also view the adapter instance within the collector. The collection status page also provides recommendations in case of any issues caused by the collection mechanism.

The collection status page summarises the number of objects and metrics that are getting collected by the VMware Cloud Foundation OperationsVMware Aria Operations collectors and adapter instances. You can view the status of each collector and adapter instance and find issues, if any. If there are anomalies, the collection status page provides recommendations to resolve those issues.

1. To view your collection status, from the left menu, click **Administration › Control Panel**, and then click the **Collection Status** tile.

**Table 286: Collection Status Overview Options**

| Option | Description |
|---|---|
| Overview | You can view the overall number of objects and metrics that is collected by your cluster. The charts underneath the numbers represent the graphical view of the collected data. Hover over the charts to view their values. |
| Collectors | You can view the total number of collectors, including data collectors in your cluster. If any of the collectors stop working, it is reported here. |
| Adapter Instances | You can view the total number of adapter instances that are receiving data. If any of the adapter instances do not receive data, it is reported here. |

2. Click the **Show Issues Only** check box to filter collectors and adapter instances that have issues.
3. By default, the collection cycle is displayed in the **Topology** structure. Using the topology view, you can clearly view the flow of data from the adapter instances to the collectors, and from the collectors into the cluster.
   Alternatively, you can click the **List View** icon to view the collection cycle in a tabular structure.

**Table 287: Collection Status Topology View**

| Option | Description |
|---|---|
| Cluster | The VMware Aria OperationsVMware Cloud Foundation Operations cluster collects data using its collectors. Hover over the cluster to view the name and type of the cluster. |
| Collectors | All the collectors that are part of a cluster are listed. You can view the number of objects and metrics being collected. Hover over the collector to view the name and type of the collector.<br>Click the **Expand** icon to view the adapter instances. |
| Adapter Instances | All the adapter instances that are part of a collector are listed. You can view the number of objects and metrics that are being collected. Hover over the adapter instance to view the name and type of the adapter. |

4. The collection status of each of the instances are colour coded.

**Table 288: Collection Status of Adapter Instances**

| Colour | Status | Description |
|---|---|---|
| Green | Collecting. | Resource is receiving data. |
| Grey | Stopped | Resource manually stopped by user. |
| Yellow | Warning | Resource is receiving data but has a problem. You can view the warning message and recommendation on how to resolve it. |
| Red | Failed | Resource fails to collect data due to some problem. |

5. Click the collector to view the details.

| Option | Description |
|---|---|
| Details | |
| Name | Name of the collector. |
| IP Address | Internet protocol (IP) address of the collector. |
| Status | Status of the collector.<br>**NOTE**<br>If any of the adapter instances within the collector has an anomaly, the status will reflect as warning. |
| Uptime | Total time elapsed since the collector started receiving data. |
| Creation Date | The day the collector was created. |
| Last HeartBeat | The last heart beat from the collector in the defined interval. |
| Version | The version of the collector. |
| Adapters | Total number of adapter instances in the collector. |
| Performance Details | |
| CPU | The average percentage of CPU used by the collector. |
| Memory | The percentage of memory used by the collector. |
| Data Collection Details | |
| Objects | Number of objects collected. |
| Metrics | Number of metrics collected. |
| Started Adapters | Number of adapter instances receiving data. |
| Threads | Number of threads collector service that is being used. |

6. Click the adapter instance to view its details.

| Option | Description |
|---|---|
| Details | |
| Name | Name of the adapter instance. |
| Status | Status of the adapter instance.<br>**NOTE**<br>If the adapter instance has a problem, the status will reflect as warning and also display the recommendation to resolve it. |
| Data Collection Details | |
| Objects | Number of objects collected. |
| Metrics | Number of metrics collected. |
| Events | Number of events collected. |
| New Objects | Determines whether new objects are collected. |
| New Metrics | Determines whether new metrics are collected. |
| New Properties | Determines whether new properties are collected. |
| Property Value Changes | Determines if the property values have changed. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Relationship Updates | Determines if there are any changes in the relationship. |
| Elapsed Collection Time | Duration of the last collection cycle. |

## Managing Content

As a VMware Cloud Foundation Operations administrator, you can take regular backups of your custom and the out-of-the-box content to manage your operational or regulatory needs. If there is a wrong edit or if the need to recover data arises, then you can use the recent backup to restore the content or import the content to a different setup. By taking regular backups, you can also upgrade the VMware Cloud Foundation Operations to the latest build without losing or overriding the custom content.

As a VMware Cloud Foundation Operations administrator, you can migrate all your content from VMware Cloud Foundation Operations to VMware Cloud Foundation Operations. You can take regular backups of your custom and the out-of-the-box content to manage your operational or regulatory needs. If there is a wrong edit or if the need to recover data arises, then you can use the recent backup to restore your content.

> **NOTE**
> Any user with the "Content Management" permission can export the content. However, only a admin user or the user assigned to the Administrator role has the privilege to export all the content, including the content owned by other users, for example, custom dashboards.

### Creating a Backup

You can create regular backups of your custom and the out-of-the-box content in VMware Aria OperationsVMware Cloud Foundation Operations. You can use this backup to restore your content or export the content while setting up another environment.

You can select the content or the configuration type that you want to export. Some content types have dependencies on other types. The dependencies, when not selected, will not be exported. Select all the required types to ensure that are no missing items when importing.

1. From the left menu, click **Administration › Control Panel**, and then click the **Content Management** tile.
2. In the **Export** tab, select the content and configurations that you want to export. You can take a backup of the following content types and configurations available in VMware Aria OperationsVMware Cloud Foundation Operations.

   > **NOTE**
   > Click **Include out-of-the-box content** under the Content section to include the out-of-the-box content while exporting. Click **Select All** under both Content and Configuration sections to select all the content and configurations respectively, while exporting.

| Content | |
|---|---|
| • Dashboards<br>• Views<br>• Reports<br>• Report Schedules | • Recommendations<br>• Notifications<br>• Payload Templates<br>• Policies |

*Table continued on next page*

*Continued from previous page*

| | |
|---|---|
| • Configuration Files<br>• Alert Definitions<br>• Symptom Definitions | • Custom Groups<br>• Custom Metric Groups<br>• Super Metrics<br>• Compliance Scorecards |
| Configuration | |
| • Authentication Sources<br>• Users<br>• User Groups<br>• User Roles<br>• Integration Accounts<br>• Http Proxies<br>• Outbound Settings | • Cost Drivers<br>• Custom Services<br>• Service Based Applications<br>• Rule Based Applications<br>• Application Definition Assignments<br>• Custom Profiles<br>• Global Settings |
| • User Groups<br>• User Roles<br>• Integration Accounts<br>• Outbound Settings<br>• Cost Drivers<br>• Custom Services | • Service Based Applications<br>• Rule Based Applications<br>• Application Definition Assignments<br>• Custom Profiles<br>• Global Settings |

**NOTE**
VMware Cloud Foundation Operations SaaS does not support the export/import of Authentication Sources, Users, and Http Proxies.

3. Click **Export** to create a backup.

   **NOTE**
   For configurations such as Integration Accounts, Http Proxies, Outbound Settings, Users, and Authentication Sources that have sensitive information, you must set up a new password to export data. The password should be at least 14 characters long.

   The system compresses the content into one ZIP file. Once the export is complete, the Download ZIP file link is available in the Export tab.

4. Click the **Download ZIP file** link to download the backup content.
   You can use the downloaded content to restore your content or export it to a different setup.

   **NOTE**
   The ZIP file generated by a certain user will not overwrite that of another user, that is, if User A generates the ZIP file after User B, the latter's ZIP file will not be overwritten.

### Importing Content

You can take regular backups of your custom and the out-of-the-box content and import it to a different environment.

- Ensure that you have downloaded the backup ZIP file. For details, see Creating a Backup.
- Ensure that all the users who own the custom dashboards or report schedules are present in the destination setup so that the custom content is assigned to the respective owners when the content is imported. Otherwise, the custom content of the owners who are not present in the destination setup will be skipped while importing the content.

**NOTE**
Items owned by a particular user (for example, dashboards and report schedules), except those owned by the admin user, will be skipped while importing the content.

**NOTE**
Only one export or import operation can take place at a time.

1. From the left menu, click **Administration › Control Panel**, and then click the **Content Management** tile.
2. Click the **Import** tab and then, click **Browse** to select the downloaded ZIP file with the exported content.
   The data included in the ZIP file is displayed in the 'Data Available' table.
3. If there is a conflict while importing the content, you can select to either **Overwrite existing content** or **Skip item(s)**.
   The import report with the timestamp is displayed after the import operation is complete. You can view this information under the **Results** section in the same page.
4. For content types with sensitive information, enter the password that you had set while exporting the content.
5. Click **Import**.
   After the import is completed, the content is available in the destination setup.

   **NOTE**
   VMware Cloud Foundation Operations SaaS does not support the export/import of Authentication Sources, Users, and Http Proxies.

#### Best Practices for Migrating Content

Follow the below practices to ensure that your content is successfully migrated.

- Use the `admin` user account or a user assigned to the Administrator role to export all of the content, including other users' custom content, such as dashboards and report schedules.
- Before importing the content, ensure that the Management Packs to which the content is related is installed on the destination setup.
- Use a user from the CSP Admin group to import all content.

## Managing Orphaned and Unassigned Content

Dashboards, report schedules, and credentials created by deleted users or that are unassigned in VMware Aria OperationsVMware Cloud Foundation Operations are stored in the Orphaned and Unassigned Content page. As an administrator, you are the new owner of these dashboards, reports schedules, and credentials. You can transfer the ownership of dashboards and report schedules and assign the credentials to new users or to yourself.

**From Where You Can Transfer Ownership of Dashboards, Report Schedules, and Credentials**

In the menu, click **Administration › Control Panel**, and then click the **Orphaned and Unassigned** tile.

**Orphaned and Unassigned Page**

You can view a list of all deleted users or unassigned users from the **Deleted Users** panel in the left pane of the **Orphaned and Unassigned** page. Based on your selection in the **Deleted Users** panel, the dashboards, report schedules, and credentials for the deleted user are displayed under the **Dashboard**, **Report Schedules**, **Credentials** tabs in the **Orphaned and Unassigned** page.

As an admin user, you can take ownership, assign ownership, or discard orphaned dashboards, report schedules and credentials, from the **Actions** menu in the **Dashboards**, **Report Schedules**, and **Credentials** tabs. Enter the name or part of the name of a dashboard, report schedule, or credential in the **Filter** option and click **Enter**. The relevant dashboard, report schedule, or credential is displayed.

**Table 289: Actions Menu Options**

| Actions | Options |
|---|---|
| Take Ownership | You can take ownership of the selected dashboard, report schedule, or credential. |
| Assign Ownership | You can assign a new owner for the selected dashboards or report schedules. You can select a target user from the **Transfer Dashboards/Report Schedule/Credential** dialog box. |
| Discard | You can permanently delete the dashboard, report schedule, and unassigned credentials that are not in use in any adapter instance. You cannot delete an active credential that is being used by an adapter instance even if it is unassigned. For more information on credentials, see Configuring Credentials in Integrations. |

## Adding Physical Data Centers in VMware Aria OperationsVMware Cloud Foundation Operations

In VMware Aria OperationsVMware Cloud Foundation Operations, you can create and manage physical data centers and associate cloud accounts with them.

1. From the left menu, click **Administration › Control Panel**, and then click the **Physical Data Centers** tile.
2. On the Physical Data Centers page, click **Add**.
   Add Physical Data Centers page opens.
3. Enter a **Name** for the physical data center.
4. Select a location on the map.

   You can use the search bar to find a specific location. You can also zoom in or zoom out for a better view of the world map and find a suitable location.

5. Double-click on the selected location to create the data center.
   The location name and coordinates are reflected in the **Location** field.
6. From the **Associate Cloud Account(s) with this Physical Data Center** section, select the cloud account(s) you want to assign to this physical data center.

   **NOTE**
   The Associate Cloud Account(s) with this Physical Data Center section only lists the configured vCenter and VMware Cloud Foundation cloud accounts. You can assign multiple cloud accounts to a physical data center.

7. Click **Save**.
   The physical data center is created.
8. View the new physical data center on the Physical Data Center page.

| Option | Description |
|---|---|
| Name | Name of the physical data center you entered at the time of creation. |
| Number of associated cloud accounts | Displays the total number of cloud accounts associated with the physical data center. |
| Filter | Filters the list of physical data centers based on Name. |

9. Click the vertical ellipsis and select **Edit** to edit the physical data center.
10. Click the vertical ellipsis or the horizontal ellipsis and then select **Delete** to remove a physical data center.

## VMware Aria OperationsVMware Cloud Foundation Operations Maintenance Schedules

Maintenance Schedules
Maintenance schedules identify objects that are in maintenance mode at specific times, which prevents VMware Aria OperationsVMware Cloud Foundation Operations from showing misleading data based on those objects being offline or in other unusual states because of maintenance.

Many objects in the enterprise might be intentionally taken offline. For example, a server might be deactivated to update software. If VMware Aria OperationsVMware Cloud Foundation Operations collects metrics when an object is offline, it might generate incorrect anomalies and alerts that affect the data for setting dynamic thresholds for the object attributes. When an object is identified as being in maintenance mode, VMware Aria OperationsVMware Cloud Foundation Operations does not collect metrics from the object or generate anomalies or alerts for it. In addition, VMware Aria OperationsVMware Cloud Foundation Operations cancels any active symptoms and alerts for the object.

If an object undergoes maintenance at fixed intervals, you can create a maintenance schedule and assign it to the object. For example, you can put an object in maintenance mode from midnight until 3 a.m. each Tuesday night. You can also manually put an object in maintenance mode, either indefinitely or for a specified period of time. These methods are not mutually exclusive. You can manually put an object in maintenance mode, or take it out of maintenance mode, even it if has an assigned maintenance schedule. For more information, see Manage Maintenance Schedules for Your Object Workspace.

### How Maintenance Schedules Work

Maintenance schedules require that you select the days and time-of-day when updates or other object maintenance occurs. Note that creating a maintenance schedule does not activate the schedule. A maintenance schedule must be part of a policy before the schedule can take effect. For more information, see Maintenance Schedule Details.

### Where You Find the Maintenance Schedules

From the left menu, click **Operations › Configurations**, and then click the **Maintainance Schedules** tile..

Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

**Table 290: Maintenance Schedule Toolbar Options**

| Option | Description |
|---|---|
| Add | Open a window in which you can select the maintenance schedule settings for a new schedule. |
| Edit | Change the maintenance schedule settings for an existing schedule. |
| Delete | Remove the selected maintenance schedule. |

## Manage Maintenance Schedules

Add or edit a maintenance schedule to take an object offline. VMware Aria OperationsVMware Cloud Foundation Operations does not collect data from an object that is offline.

**Where You Find Manage Maintenance Schedules**

1. From the left menu, click **Operations › Configurations**, and then click the **Maintainance Schedules** tile.
2. Click **Add** or click the **Vertical Ellipses** to edit, or remove items.

**Table 291: Manage Maintenance Schedule Add or Edit Options**

| Option | Description |
|---|---|
| Schedule Name | Name that describes the maintenance schedule |
| Time Zone | Time zone in which you are currently located |
| Days | Number of days the maintenance period covers |
| Recurrence | Specify a maintenance schedule to run over a selected period<br>• Once<br>• Daily<br>• Weekly<br>• Monthly |
| Expire after | The number of times the schedule is run |
| Expire on | The date upon which the schedule stops running |

## Create a VMware Aria OperationsVMware Cloud Foundation Operations Support Bundle

Create a Support Bundle
You create a VMware Aria OperationsVMware Cloud Foundation Operations support bundle to gather log and configuration files for analysis when troubleshooting a VMware Aria OperationsVMware Cloud Foundation Operations issue.

When you create a support bundle, VMware Aria OperationsVMware Cloud Foundation Operations gathers files from cluster nodes into ZIP files for convenience.

1. In the menu, click **Administration › Control Panel**, and then click the **Support Bundles** tile.
2. From the toolbar, click the **Create a Support Bundle** icon.
3. Select the option to create a **Light** or **Full support bundle**.
4. Select the cluster nodes that need to be evaluated for support.

   Only logs from the selected nodes are included in the support bundle.

5. Click **OK**, and click **OK** to confirm support bundle creation.

   Depending on the size of the logs and number of nodes, it might take time for VMware Cloud Foundation Operations to create the support bundle.

Use the toolbar to download the support bundle ZIP files for analysis. For security, VMware Aria OperationsVMware Cloud Foundation Operations prompts you for credentials when you download a support bundle.

You can review the log files for error messages or, if you need troubleshooting assistance, send the diagnostic data to VMware Technical Support. When you resolve or close the issue, use the toolbar to delete the outdated support bundle to save disk space.

### VMware Aria OperationsVMware Cloud Foundation Operations Support Bundles

Support Bundles
VMware Aria OperationsVMware Cloud Foundation Operations support bundles contain log and configuration files that help troubleshoot a VMware Aria OperationsVMware Cloud Foundation Operations issue.

#### How Support Bundles Work

Support bundles require that you select nodes or the entire cluster, and the level of logging that you want to collect. After VMware Aria OperationsVMware Cloud Foundation Operations creates the support bundle, you download it in ZIP format for analysis.

#### Where You Find Support Bundles

In the menu, click **Administration › Control Panel**, and then click the **Support Bundles** tile.

#### Support Bundle Options

The options include toolbar and data grid options.

You can click **Add** or click the **Hortizontal Ellipses** to delete, download, or reload support bundles.

**Table 292: Support Bundle Toolbar Options**

| Option | Description |
|---|---|
| Add | Open a dialog box that guides you through the process of creating a support bundle. |
| Delete | Remove the selected support bundle. |
| Download | Download the support bundle in ZIP format. |
| Reload Support Bundles | Refresh the list of support bundles. |

Use the data grid options to view item details.

**Table 293: Support Bundle Data Grid Options**

| Option | Description |
|---|---|
| Bundle | System-generated identifier for the support bundle. |
| Bundle Type | • Light. Include 24 hours of logs.<br>• Full. Include all available logs and configuration files. |
| Date and Time Created | Time when support bundle creation began. |
| Status | Progress of support bundle creation. |

## VMware Aria OperationsVMware Cloud Foundation Operations Dynamic Thresholds

Dynamic Thresholds

A threshold marks the boundary between normal and abnormal behavior for a metric. In addition to fixed thresholds, VMware Aria OperationsVMware Cloud Foundation Operations supports dynamic thresholds for a metric, calculated based on historical and incoming data.

### How Dynamic Thresholds Work

By default, dynamic thresholds are refreshed on a regular schedule, but you can recalculate dynamic thresholds outside of the schedule if you want to capture the most recent data.

### Where You Find Dynamic Thresholds

In the menu, click **Administration › Global Settings**, and then click the **Dynamic Thresholds** tile.

### Dynamic Threshold Options

The dynamic threshold feature includes options to activate or deactivate the calculation process and to review associated values.

**Table 294: Dynamic Threshold Options**

| Option | Description |
|---|---|
| Activate | The dynamic threshold calculation process ia activated by default. |
| Deactivate | Stop the dynamic threshold calculation currently in progress. |
| Calculation progress | Percentage completion of the current dynamic threshold calculation. |
| Host time to calculate normal behavior of all objects | Timestamps and metric counts associated with the last dynamic threshold calculation, and the time for the next scheduled calculation. |

## VMware Aria OperationsVMware Cloud Foundation Operations Adapter Redescribe

Adapter Redescribe
When VMware Aria OperationsVMware Cloud Foundation Operations redescribes an adapter, VMware Aria OperationsVMware Cloud Foundation Operations finds the adapter files, gathers information about the abilities of the adapter, and updates the user interface with information about the adapter.

### How Adapter Redescribe Works

After installing or updating an adapter, capture the adapter information by having VMware Aria OperationsVMware Cloud Foundation Operations redescribe its adapters.

### Where You Find Adapter Redescribe

In the menu, click **Administration › Control Panel**, and then click the **Redescribe** tile.

### Adapter Redescribe Options

The feature includes an option to start the adapter describe process.

**Table 295: Adapter Redescribe Options**

| Option | Description |
|---|---|
| Redescribe | Start the adapter describe process. |

VMware Aria OperationsVMware Cloud Foundation Operations provides adapter-specific details from the redescribe process.

**Table 296: Adapter Redescribe Details**

| Option | Description |
|---|---|
| Name | Adapter to which the redescribe process applies. |
| Status | Success, failure, or other condition related to the last redescribe process. |
| Describe Version | Version of `describe.xml` against which the last redescribe process ran. |
| Adapter Version | Version of the adapter against which the last redescribe process ran. |
| Message | Additional details about the last redescribe process. |

## Customizing Icons

Every object or adapter in your environment has an icon representation. You can customize how the icon appears.

VMware Aria OperationsVMware Cloud Foundation Operations assigns a default icon to each object type and adapter type . Taken collectively, object types and adapter types are known as objects in your environment. Icons represent objects in the UI and help you to identify the type of object. For example, in the Topology Graph widget on a dashboard, labeled icons show how objects are connected to one other. You can quickly identify the type of object from the icon.

If you want to differentiate objects, you can change the icon. For example, a virtual machine icon is generic. If you want to pictorially distinguish the data that a vSphere virtual machine provides from the data that a Hypervisor virtual machine provides, you can assign a different icon to each.

### Customize an Object Type Icon

You can use the default icons that VMware Aria OperationsVMware Cloud Foundation Operations provides, or you can upload your own graphics file for an object type. When you change an icon, your changes take effect for all users.

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

1. From the left menu, click **Administration › Control Panel**, and then click the **Icons** tile.
2. Click the **Object Type Icons** tab.
3. Assign the Object Type icon.
   a) Select the object type in the list with the icon to change.

      By default, object types for all adapter types are listed. To limit the selection to the object types that are valid for a single adapter type, select the adapter type from the drop-down menu.

   b) Click the **Upload** icon.
   c) Browse to and select the file to use and click **Done**.
4. To return to the default icon, select the object type and click the **Assign Default Icons** icon.
   The original default icon appears.

### Object Type Icons Tab

VMware Aria OperationsVMware Cloud Foundation Operations obtains data from different sources. Data sources are classified by the type of object or object type. In UI locations where metric data appears for objects, VMware Aria

OperationsVMware Cloud Foundation Operations includes an icon to show the object type. To graphically distinguish the different types of objects, you can customize the icon.

**Where You Customize Object Type Icons**

From the left menu, click **Administration › Control Panel**, and then click the **Icons** tile. Click the **Object Type Icons** tab.

**Table 297: Object Type Icons Options**

| Option | Description |
|---|---|
| Adapter Type | Icons for all adapters are listed by default. To list a subset of the object types that are valid for one type of adapter, select the adapter type. |
| Toolbar options | Manages the selected icon.<br>• **Upload** uploads a PNG file to uniquely identify the object type.<br>• **Assign Default icons** returns the selection to the original icon. |
| Search | Search for objects with a particular name to narrow the selection of object types displayed. |
| Object Type | Name of the type of object. |
| Icon | Pictorial representation of the type of object. |

**Customize an Adapter Type Icon**

You can use the default icons that VMware Aria OperationsVMware Cloud Foundation Operations provides, or you can upload your own graphics file for an adapter type. When you change an icon, your changes take effect for all users.

If you plan to use your own icon files, verify that each image is in PNG format and has the same height and width. For best results, use a 256x256 pixel image size.

1. From the left menu, click **Administration › Control Panel**, and then click the **Icons** tile.
2. Click the **Adapter Type Icons** tab.
3. Assign the Adapter Type icon.
   a) Select the adapter type in the list with the icon to change.
   b) Click the **Upload** icon.
   c) Browse to and select the file to use and click **Done**.
4. To return to the default icon, select the adapter type and click the **Assign Default Icons** icon.
   The original default icon appears.

**Adapter Type Icons Tab**

Adapters collect and provide data to VMware Aria OperationsVMware Cloud Foundation Operations. Adapters are classified by the type of adapter or adapter kind. To graphically distinguish the different types of adapters, you can customize the icon.

**Where You Customize Adapter Type Icons**

From the left menu, click **Administration › Control Panel**, and then click the **Icons** tile. Click the **Adapter Type Icons** tab.

**Table 298: Adapter Type Icons Options**

| Option | Description |
|---|---|
| Toolbar options | Manages the selected icon.<br>• **Upload** uploads a PNG file to uniquely identify the adapter type.<br>• **Assign Default icons** returns the selection to the original icon. |
| Name | Name of the type of adapter. |
| Icon | Pictorial representation of the type of adapter. |

## Allocate More Virtual Memory to VMware Aria OperationsVMware Cloud Foundation Operations

Allocate More Virtual Memory
You might need to add virtual memory to keep the VMware Aria OperationsVMware Cloud Foundation Operations process running.

When the VMware Aria OperationsVMware Cloud Foundation Operations virtual machine requests more memory than is available, the Linux kernel might kill the `vcops-analytics` process, and the product might become unresponsive. If that happens, use the reservation feature in vSphere to specify the guaranteed minimum memory allocation for VMware Aria OperationsVMware Cloud Foundation Operations virtual machines.

1. In the vSphere Client inventory, right-click the VMware Aria OperationsVMware Cloud Foundation Operations virtual machine and select **Edit Settings**.
2. Click the **Resources** tab, and select **Memory**.
3. Use the **Reservation** option to allocate more memory.

## About the VMware Cloud Foundation OperationsVMware Aria Operations Administration Interface

About the Administration Interface
The VMware Cloud Foundation OperationsVMware Aria Operations administration interface provides access to selected maintenance functions beyond what the product interface supports.

Use the VMware Cloud Foundation OperationsVMware Aria Operations administration interface instead of the product interface under the following conditions. You can access the administration interface login page from any node in the VMware Cloud Foundation OperationsVMware Aria Operations analytics cluster by appending `/admin` to the node IP address or FQDN when you enter the URL in your browser.

- Activate or deactivate high availability (HA).
- Upload and install VMware Cloud Foundation OperationsVMware Aria Operations software update PAK files.
- The product interface is inaccessible, and you must correct the problem by bringing nodes online, or by restarting nodes or the cluster.
- VMware Cloud Foundation OperationsVMware Aria Operations needs to be restarted for any reason.

There is some overlap between the administration interface and product interface in terms of access to logs, support bundles, and some of the node maintenance activities that do not involve restarting the cluster, such as adding nodes.

## Configuring the Global Network Time Protocol (NTP) Settings

The Network Time Protocol (NTP) setting is used to synchronize the time of your VMware Cloud Foundation OperationsVMware Aria Operations cluster nodes. Configure the NTP setting in VMware Cloud Foundation

OperationsVMware Aria Operations to ensure that all the cluster nodes follow the same time protocol. Following the same time protocol allows your cluster's services to run smoothly.

The primary node serves as the NTP server by default and the time of all nodes are synchronized to the primary node. You can continue using the primary node as the NTP server and add another NTP server to serve as a failover if the primary node fails. In case you do not want to use the primary node, you can add the NTP server you wish to use.

1. Log in to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://*primary-node-name-or-ip-address*/admin.
2. Click the **NTP Settings** icon.
   The Global Network Time Protocol settings wizard opens. The primary node is listed as the existing NTP server.

   **NOTE**
   You must configure at least two NTP servers to have a back-up if one server fails.

3. Add a secondary NTP server. Enter the IP of FDQN of the NTP server you wish to use in **NTP Server Address** field.
4. Click **Add**.
   The new NTP server is displayed along with the primary node.
5. Ensure that the status of the NTP server is green which means its reachable and working.
6. To delete a NTP server, click **Remove** next to the NTP server status.

## VMware Cloud Foundation OperationsVMware Aria Operations Cluster Management

Cluster Status and Management
VMware Cloud Foundation OperationsVMware Aria Operations includes a central page where you can monitor and manage the nodes in your VMware Cloud Foundation OperationsVMware Aria Operations cluster and the adapters that are installed on the nodes.

### How Cluster Management Works

You can view and change the online or offline state of the overall VMware Cloud Foundation OperationsVMware Aria Operations cluster or the individual nodes. In addition, you can activate or deactivate high availability (HA) and view statistics related to the adapters that are installed on the nodes.

### Where You Find Cluster Management

Log in to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://*master-node-name-or-ip-address*/admin.

### Cluster Management Options

The options include cluster-level monitoring and management features.

**Table 299: Initial Setup Status Details**

| Option | Description |
|---|---|
| Cluster Status | Displays the online, offline, or unknown state of the VMware Cloud Foundation OperationsVMware Aria Operations cluster and provides an option to take the cluster online or offline.<br>If a cluster fails to go offline, click the **Force Take Offline** button to take the cluster offline. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | **NOTE**<br>The Force Take Offline button appears only when the Bring Cluster offline operation fails.<br><br>You can select to display the reason for taking the cluster offline. Select the **Show reason on maintenance page** check box in the **Take Cluster Offline** dialog box. When you log in to VMware Cloud Foundation OperationsVMware Aria Operations when the cluster is offline, the reason for taking the cluster offline is displayed. |
| High Availability | Indicates whether HA is activated, deactivated, or degraded and provides an option to change that setting. |
| Continuous Availability | Indicates whether CA is activated, deactivated, or degraded and provides an option to change that setting. |

VMware Cloud Foundation OperationsVMware Aria Operations provides node-level information as well as a toolbar for taking nodes online or offline.

**Table 300: Nodes in the VMware Cloud Foundation OperationsVMware Aria Operations Cluster**

| Option | Description |
|---|---|
| Generate Passphrase | Generate a passphrase that can be used instead of the administrator credentials to add a node to this cluster. |
| Add New Node | Add a new node to the this cluster. You cannot add a witness node. |
| Take Node Online/Offline | You can select the required node and bring it online or offline. You are required to understand the risk involved and provide a valid reason for the action performed when you bring a node online or offline. |
| Remove Node | Remove node from the cluster without any loss of collected data. Data nodes must be removed by shrinking. |
| Reload Nodes | Reload data in the screen. |
| Shrink Cluster | This option provides a mechanism to remove a node without having to lose any data. The shrink cluster removes nodes by migrating data from one node to any other node. All the historical data is either moved to the primary node or any other node, which has sufficient disk space.<br><br>If HA is activated and you have selected the replica node for removal, then you are asked to select another replica node. VMware Cloud Foundation OperationsVMware Aria Operations provides a list of nodes that be a possible candidate to become a replica node.<br><br>VMware Cloud Foundation OperationsVMware Aria Operations stops collecting data from the removed nodes. However, the data that is available in the removed node is migrated to an existing node. Once the migration is |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | complete, then the removed nodes are deleted with the cluster state as offline.<br>**NOTE**<br>VMware Cloud Foundation OperationsVMware Aria Operations cannot move pinned adapters. The adapter instances which were pinned on removed nodes do not move to another collector automatically. You must change the collector before starting the shrink cluster process. |
| Replace Node | You cannot remove one node in a CA cluster, because the nodes are in pairs. Use the replace node option to replace one of the nodes in the CA cluster. The following rules apply:<br>• Replace master node - not allowed<br>• Replace data node -<br>– Single data node - allowed<br>– Data nodes from different FDs - allowed<br>– Pair of data nodes - not allowed<br>– Data nodes from the same FD - not allowed<br>• Replace witness node - allowed |

**Table 301: Nodes in the VMware Cloud Foundation OperationsVMware Aria Operations Cluster**

| Option | Description |
|---|---|
| Node Name | Machine name of the node.<br>The node that you are logged into displays a dot next to the name. |
| Node Address | Internet protocol (IP) address of the node. Primary and replica nodes require static IP addresses. Data nodes may use DHCP or static IP. |
| Cluster Role | Type of VMware Cloud Foundation OperationsVMware Aria Operations node: primary, data, or replica. |
| State | Powered on, powered off, unknown, or other condition of the node. |
| Status | Online, offline, unknown, or other condition of the node. |
| Objects | Total environment objects that the node currently monitors. |
| Metrics | Total metrics that the node has collected since being added to the cluster. |
| Build | VMware Cloud Foundation OperationsVMware Aria Operations software build number installed on the node. |
| Version | VMware Cloud Foundation OperationsVMware Aria Operations software version installed on the node. |
| Deployment Type | Type of machine on which the node is running: vApp |
| SSH Status | Activate or deactivate the SSH Status. |

In addition, there are adapter statistics for the selected node.

**Table 302: Adapters on Server**

| Option | Description |
|---|---|
| Name | Name that the installing user gave to the adapter. |
| Status | Indication of whether the adapter is collecting data or not. |
| Objects | Total environment objects that the adapter currently monitors. |
| Metrics | Total metrics that the adapter has collected since being installed on the node. |
| Last Collection Time | Date and time of the most recent data collection by the adapter. |
| Added On | Date and time when the adapter was installed on the node. |

## Monitoring the Health of Cloud Proxies from the Admin UI

After you configure your cloud proxy, you can view the status, health, and upgrade history of your cloud proxy in the VMware Aria Operations administration interface.

1. Log in to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://*primary-node-name-or-ip-address*/admin.
2. Click **Cloud Proxies**.

   **Table 303: Cloud Proxies Page Options**

   | Option | Description |
   |---|---|
   | IP Address | IP address of the cloud proxy. |
   | Name | Name of the cloud proxy. |
   | Network Proxy Configuration | Determines whether the network proxy setting is configured or not. |
   | Health Status | Determines the health of the cloud proxy. |
   | Upgrade Status | Determines whether the upgrade is complete, in progress, or failed. |
   | Last Upgrade Time | Determines when was the last upgrade done. |
   | Version | Version number of the cloud proxy. |

3. Click the **Expand** icon to view the upgrade history.

   **Table 304: Upgrade History Options**

   | Option | Description |
   |---|---|
   | ID | The conventional name used to identify the PAK file. It is usually the name of the PAK file and its version numbers joined together without extensions. For example, VMware Aria-Operations-Cloud-Proxy-84045207710. |
   | Type | The type of upgrade used for the cloud proxy. The cloud proxy can be upgraded either automatically or manually using the command-line interface. For more information, see the topic called Using the Cloud Proxy Command- |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | Line Interface in the *Getting Started with VMware Cloud Foundation OperationsVMware Aria Operations* guide. |
| Start Time | Timestamp when the upgrade started. |
| End Time | Timestamp when the upgrade ended. |
| Upgrade Status | Determines whether the upgrade is complete, in progress, or failed.<br>**NOTE**<br>For a detailed view of the upgrade status, click the row to open the **Info** pop-up. If the cloud proxy upgrade fails due to low disk space, the `FAIL_NO_SPACE_FOR_PAK_DOWNLOAD` message appears. |
| Version | Version number of the cloud proxy PAK file. |

## VMware Cloud Foundation OperationsVMware Aria Operations Logs for Admin UI

Logs
For troubleshooting in the Admin UI, the product provides an expandable tree of VMware Cloud Foundation OperationsVMware Aria Operations log files that you can browse and load for review.

### How VMware Cloud Foundation OperationsVMware Aria Operations Logs Work

VMware Cloud Foundation OperationsVMware Aria Operations logs are categorized by cluster node, and functional area or log type.

### Where You Find VMware Cloud Foundation OperationsVMware Aria Operations Logs

Log in to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://*master-node-name-or-ip-address*/admin and then click **Support › Logs**.

### Log Viewer Options

Use the toolbar options to control the tree of items and the viewer.

**Table 305: Log Viewer Toolbar Options**

| Option | Description |
|---|---|
| Starting Line | Specifies the starting line of the file to be displayed.<br>Note: 0 is for the first line. -1 or no value indicates that the file has to be displayed from the end. |
| Number of Lines | Specifies the number of lines to be displayed from the file. For example:<br>If you want to see the first 10 lines of the required text, specify the number of lines as 10 and the starting line as 0. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Word Wrap | If you select this option, the extra part of the line that does not fit on the screen is moved to the next line. If you do not select this option, a scroll bar is provided to see the complete line. |

## VMware Cloud Foundation OperationsVMware Aria Operations Support Bundles

Support Bundles
VMware Cloud Foundation OperationsVMware Aria Operations support bundles contain log and configuration files that help troubleshoot a VMware Cloud Foundation OperationsVMware Aria Operations issue.

### How Support Bundles Work

Support bundles require that you select nodes or the entire cluster, and the level of logging that you want to collect. After VMware Cloud Foundation OperationsVMware Aria Operations creates the support bundle, you download it in ZIP format for analysis.

### Where You Find Support Bundles

Log in to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://*master-node-name-or-ip-address*/admin and then click **Support › Support Bundles**.

### Support Bundle Options

The options include toolbar and data grid options.

Use the toolbar options to add, download, or remove items.

**Table 306: Support Bundle Toolbar Options**

| Option | Description |
|---|---|
| Add | Open a dialog box that guides you through the process of creating a support bundle. |
| Delete | Remove the selected support bundle. |
| Download | Download the support bundle in ZIP format. |
| Reload | Refresh the list of support bundles. |

Use the data grid options to view item details.

**Table 307: Support Bundle Data Grid Options**

| Option | Description |
|---|---|
| Bundle | System-generated identifier for the support bundle. |
| Bundle Type | • Light. Include 24 hours of logs.<br>• Full. Include all available logs and configuration files. |
| Date and Time Created | Time when support bundle creation began. |
| Status | Progress of support bundle creation. |
| File Size | The size of the support bundles. |

## Support Bundles (Cloud Proxy)

VMware Cloud Foundation OperationsVMware Aria Operations support bundles contain log and configuration files that help troubleshoot a VMware Cloud Foundation OperationsVMware Aria Operationscloud proxy issue.

Use the Support Bundles (Cloud Proxy) page to create the support bundle on cloud proxy.

Log in to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://*master-node-name-or-ip-address*/admin and then click **Support › Support Bundles (Cloud Proxy)**.

### Support Bundle (Cloud Proxy) Options

The options include toolbar and data grid options.

Use the toolbar options to add, download, or remove items.

**Table 308: Support Bundle (Cloud Proxy) Toolbar Options**

| Option | Description |
|---|---|
| Add | Open a dialog box that guides you through the process of creating a support bundle on cloud proxy.<br>Select the cloud proxy and then click **OK** to create a support bundle on the selected cloud proxy.<br><br>The support bundle is created under the following directory /storage/core/vmware-vrops-cprc/support. |
| Delete | Remove the selected support bundle. |
| Download | Download the support bundle in ZIP format. |
| Reload | Refresh the list of support bundles. |

Use the data grid options to view item details.

**Table 309: Support Bundle (Cloud Proxy) Data Grid Options**

| Option | Description |
|---|---|
| Bundle | System-generated identifier for the support bundle. |
| Cloud Proxy Name | The name of the cloud proxy on which the support bundle is created. |
| Date and Time Created | Time when support bundle creation began. |
| Status | Progress of support bundle creation. |
| File Size | The size of the support bundles. |

**NOTE**
Generation and download of support bundles through the Support Bundles (Cloud Proxy) page works only if the cloud proxy is connected to the cluster.

If there is a disconnect between the cloud proxy and VMware Cloud Foundation OperationsVMware Aria Operations, you can generate a support bundle on cloud proxy manually.

Open an SSH connection with the cloud proxy appliance and run the followng command:

- For 8.3 and 8.4 version: `$> cprc-cli -sb`
- For 8.5 and later versions: `$> cprc-cli -sb IS_HEAVY`, where `IS_HEAVY` should be specified as `true` or `false`

With `cprc-cli -sb true`, the support bundle is generated with journal.ctl logs. With `cprc-cli -sb false` , the support bundle is generated without journal.ctl logs. The support bundle is created in the following path: `directory / storage/db/vmware-vrops-cprc/support`.

For adding some extra files into the support bundle, you need to modify configuration file available here: `/storage/db/vmware-vrops-cprc/configuration/cprc.support.bundle.configuration` . Add the path of needed file into **`cprc.support.bundle.configuration`** in the **`files:`** section. Then generate the support bundle with the `cprc-cli -sb` or `cprc-cli -sb false` command, based on the version you are on.

For example if you want to add the `/var/log/firstboot/vcopssuitevm.log` file into the support bundle, add this path into the **`files:`** section of `/storage/db/vmware-vrops-cprc/configuration/cprc.support.bundle.configuration`. Then generate the support bundle with either the `cprc-cli -sb` or the `cprc-cli -sb false` command, based on the version you are on.

Once done, you can download or delete these support bundles from the Support Bundles (Cloud Proxy) page.

## Security Settings - Admin UI

You can activate Federal Information Processing Standards (FIPS) for VMware Aria Operations to make your environment FIPS compliant and you can also activate firewall hardening.

### Activate FIPS

You can activate FIPS in the VMware Aria Operations cluster at the time of installation or after VMware Aria Operations is up and running. Adding FIPS at installation is less intrusive because the cluster has not yet started.

If the cluster is running, to activate FIPS, you must take the cluster offline. For more information, see Cluster Management.

FIPS mode is supported in Cloud Proxy. You can continue using your cloud proxy after enabling FIPS for the VMware Aria Operations cluster.

1. In a Web browser, navigate to the master node administration interface. `https://master-node-name-or-ip-address/admin`.
2. Enter the VMware Cloud Foundation Operations administrator username of admin.
3. Enter the VMware Cloud Foundation Operations administrator password and click **Log In**.
4. Click **Administrator Settings**.

   **NOTE**
   The **Activate FIPS** button is deactivated when the cluster is running.

5. Click **Activate FIPS** after you take your cluster offline.

   **NOTE**
   Once you activate FIPS, you cannot deactivate the FIPS mode in the current setup. To revert to a FIPS deactivated setup, you must re-deploy VMware Aria Operations.

6. In the **Are you sure you want to activate FIPS** dialog box, read the note and provide your consent for enabling FIPS and then click **Yes**.

   **NOTE**
   Once you activate FIPS, the cluster restarts and is not be available during this time. The cluster nodes are rebooted and once the cluster is online, all the nodes are FIPS activated.

### Activate Firewall Hardening

Activating firewall hardening restricts network access to internal services in VMware Cloud Foundation Operations.

1. In a Web browser, navigate to the master node administration interface. `https://master-node-name-or-ip-address/admin`.
2. Enter the VMware Cloud Foundation Operations administrator username of admin.
3. Enter the VMware Cloud Foundation Operations administrator password and click **Log In**.
4. Click **Administrator Settings**, and then click **Security Settings** from the **Administrator Settings** page.
5. Click **Activate Firewall Hardening**.

## Custom VMware Aria OperationsVMware Cloud Foundation Operations Certificates

Custom Certificates
For secure VMware Aria OperationsVMware Cloud Foundation Operations operation, you might need to perform maintenance on authentication certificates.

Authentication certificates are for a secure machine-to-machine communication within VMware Aria OperationsVMware Cloud Foundation Operations itself or between VMware Aria OperationsVMware Cloud Foundation Operations and other systems.

By default, VMware Aria OperationsVMware Cloud Foundation Operations includes its own authentication certificates. The default certificates cause the browser to display a warning when you connect to the VMware Aria OperationsVMware Cloud Foundation Operations user interface.

Your site security policies might require that you use another certificate, or you might want to avoid the warnings caused by the default certificates. In either case, VMware Aria OperationsVMware Cloud Foundation Operations supports the use of your own custom certificate. You can upload your custom certificate during the initial primary node configuration or later.

### Custom VMware Aria OperationsVMware Cloud Foundation Operations Web Certificate Requirements

Custom Web Certificate Requirements
A certificate used with VMware Aria OperationsVMware Cloud Foundation Operations must conform to certain requirements. Using a custom certificate is optional and does not affect VMware Aria OperationsVMware Cloud Foundation Operations features. You can also use wildcard certificates in VMware Aria OperationsVMware Cloud Foundation Operations.

#### Requirements for Custom Certificates

Custom VMware Aria OperationsVMware Cloud Foundation Operations certificates must meet the following requirements.

- The certificate file must include the terminal (leaf) server certificate, a private key, and all issuing certificates if the certificate is signed by a chain of other certificates.
- In the file, the leaf certificate must be first in the order of certificates. After the leaf certificate, the order does not matter.
- In the file, all certificates and the private key must be in PEM format. VMware Aria OperationsVMware Cloud Foundation Operations does not support certificates in PFX, PKCS12, PKCS7, or other formats.
- In the file, all certificates and the private key must be PEM-encoded. VMware Aria OperationsVMware Cloud Foundation Operations does not support DER-encoded certificates or private keys.
  PEM-encoding is base-64 ASCII and contains legible BEGIN and END markers, while DER is a binary format. Also, file extension might not match encoding. For example, a generic `.cer` extension might be used with PEM or DER. To verify encoding format, examine a certificate file using a text editor.
- The file extension must be `.pem`.
- The private key must be generated by the RSA or DSA algorithm.
- The private key can be encrypted by a pass phrase. The generated certificate can be uploaded using the primary node configuration wizard or the administration interface.

- The REST API in this VMware Aria OperationsVMware Cloud Foundation Operations release supports private keys that are encrypted by a pass phrase.
- The VMware Aria OperationsVMware Cloud Foundation Operations certificate must have IPs and Hostnames in the Subject Alternative Name (SAN) extension.
  For example, Subject Alternative Name comprises of the DNS Name: localhost and the IP Address: 127.0.0.1.
- The VMware Aria OperationsVMware Cloud Foundation Operations Web server on all nodes have the same certificate file, so it must be valid for all nodes. One way to make the certificate valid for multiple addresses is with multiple Subject Alternative Name (SAN) entries.
- SHA1 certificates create browser compatibility issues. Therefore, ensure that all certificates that are created and being uploaded to VMware Aria OperationsVMware Cloud Foundation Operations are signed using SHA2 or newer.
- The VMware Aria OperationsVMware Cloud Foundation Operations supports custom security certificates with key length up to 8192 bits. An error is displayed when you try to upload a security certificate generated with a stronger key length beyond 8192 bits.

  **NOTE**
  Fill the certificate extension fields using the UTF-8 encoding.

vRealize Operations Manager 6.x fails to accept and apply Custom CA Certificate. For more information, see the following KB article 2046591.

**Configure a Custom Web Certificate**

You can use OpenSSL to configure an authentication certificate for use with VMware Aria OperationsVMware Cloud Foundation Operations. You must first generate a Certificate PEM for VMware Aria OperationsVMware Cloud Foundation Operations, then install the Certificate PEM in VMware Aria OperationsVMware Cloud Foundation Operations. The certificates applied through the VMware Aria OperationsVMware Cloud Foundation Operations Admin UI will be used only for securely connecting and serving the user interfaces to (external) clients. We do not update the SSL certificates used for establishing a secure connection from VMware Aria OperationsVMware Cloud Foundation Operations to other services like VMware Identity Manager, vCenter, and VMware Cloud Foundation Operations for logs.

Take your cluster offline before uploading the custom web certificate.

1. Generate a Certificate PEM file for use with VMware Aria OperationsVMware Cloud Foundation Operations.

   1. Generate a key pair by running this command:
      `openssl genrsa -out key_filename.key 2048`
   2. Use the key to generate a certificate signing request by running this command:
      `openssl req -new -key key_filename.key -out certificate_request.csr`
   3. Submit the CSR file to your Certificate Authority (CA) to obtain a signed certificate.
   4. From your Certificate Authority, download the certificate and the complete issuing chain (one or more certificates). Download them in Base64 format.
   5. Enter the command to create a single PEM file containing all certificates and the private key. In this step, the example certificate is *server_cert.cer* and the issuing chain is *cacerts.cer*.

      **NOTE**
      The order of CA's certs in the .PEM file: Cert, Private Key, Intermediate Cert and then Root Cert.
      cat *server_cert.cer key_filename.key cacerts.cer* > *multi_part.pem*

      In Windows replace cat with type.

   The finished PEM file should look similar to the following example, where the number of CERTIFICATE sections depends on the length of the issuing chain:

   ```
   -----BEGIN CERTIFICATE-----

   (Your Primary SSL certificate: your_domain_name.crt)
   ```

```
-----END CERTIFICATE-----

-----BEGIN RSA PRIVATE KEY-----

(Your Private Key: your_domain_name.key)

-----END RSA PRIVATE KEY-----

-----BEGIN CERTIFICATE-----

(Your Intermediate certificate: DigiCertCA.crt)

-----END CERTIFICATE-----

-----BEGIN CERTIFICATE-----

(Your Root certificate: TrustedRoot.crt)

-----END CERTIFICATE-----
```

2. Install a PEM in VMware Aria OperationsVMware Cloud Foundation Operations.
    1. In a Web browser, navigate to the VMware Aria OperationsVMware Cloud Foundation Operations administration interface.
       `https://vrops-node-FQDN-or-ip-address/admin`
    2. Log in with the admin user name and password.
    3. At the upper right, click the yellow **SSL Certificate** icon.
    4. In the **SSL Certificate** window, click **Install New Certificate**.
    5. Click **Browse** for certificate.
    6. Locate the certificate .pem file, and click Open to load the file in the **Certificate Information** text box. The certificate file must contain a valid private key and a valid certificate chain.
    7. Click **Install**.

### Verifying a Custom VMware Aria OperationsVMware Cloud Foundation Operations Web Certificate

Verifying a Custom Web Certificate
When you upload a custom certificate file, the VMware Aria OperationsVMware Cloud Foundation Operations interface displays summary information for all certificates in the file.

For a valid custom certificate file, you should be able to match issuer to subject, issuer to subject, back to a self-signed certificate where the issuer and subject are the same.

In the following example, `OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32` is issued by `OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32`, which is issued by `OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-ac4e-e1806f0d3f84`, which is issued by itself.

```
Thumbprint: 80:C4:84:B9:11:5B:9F:70:9F:54:99:9E:71:46:69:D3:67:31:2B:9C

Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32

Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-slice-32

Subject Alternate Name:

PublicKey Algorithm: RSA

Valid From: 2015-05-07T16:25:24.000Z
```

```
Valid To: 2020-05-06T16:25:24.000Z


Thumbprint: 72:FE:95:F2:90:7C:86:24:D9:4E:12:EC:FB:10:38:7A:DA:EC:00:3A

Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-
ac4e-e1806f0d3f84

Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-intermediate-32

Subject Alternate Name: localhost,127.0.0.1

PublicKey Algorithm: RSA

Valid From: 2015-05-07T16:25:19.000Z

Valid To: 2020-05-06T16:25:19.000Z


Thumbprint: FA:AD:FD:91:AD:E4:F1:00:EC:4A:D4:73:81:DB:B2:D1:20:35:DB:F2

Issuer Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-ad81-4a15-
ac4e-e1806f0d3f84

Subject Distinguished Name: OU=MBU,O=VMware\, Inc.,CN=vc-ops-cluster-ca_33717ac0-
ad81-4a15-ac4e-e1806f0d3f84

Subject Alternate Name: localhost,127.0.0.1

PublicKey Algorithm: RSA

Valid From: 2015-05-07T16:24:45.000Z

Valid To: 2020-05-06T16:24:45.000Z
```

### Sample Contents of Custom VMware Aria OperationsVMware Cloud Foundation Operations Web Certificates

Sample Contents of Custom Web Certificates
For troubleshooting purposes, you can open a custom certificate file in a text editor and inspect its contents.

#### PEM Format Certificate Files

A typical PEM format certificate file resembles the following sample.

```
-----BEGIN CERTIFICATE-----

MIIF1DCCBLygAwIBAgIKFYXYUwAAAAAGTANBgkqhkiG9w0BAQ0FADBhMRMwEQYK

CZImiZPyLGQBGRYDY29tMRUwEwYKCZImiZPyLGQBGRYFdm13Y3MxGDAWBgoJkiaJ

<snip>

vKStQJNr7z2+pTy92M6FgJz3y+daL+9ddbaMNp9fVXjHBoDLGGaLOvyD+KJ8+xba

aGJfGf9ELXM=

-----END CERTIFICATE-----

-----BEGIN RSA PRIVATE KEY-----
```

```
MIIEowIBAAKCAQEA4l5ffX694riI1RmdRLJwL6sOWa+Wf70HRoLtx21kZzbXbUQN
mQhTRiidJ3Ro2gRbj/btSsI+OMUzotz5VRT/yeyoTC5l2uJEapld45RroUDHQwWJ
<snip>
DAN9hQus3832xMkAuVP/jt76dHDYyviyIYbmxzMalX7LZy1MCQVg4hCH0vLsHtLh
M1rOAsz62Eht/iB61AsVCCiN3gLrX7MKsYdxZcRVruGXSIh33ynA
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIIDnTCCAoWgAwIBAgIQY+j29InmdYNCs2cK1H4kPzANBgkqhkiG9w0BAQ0FADBh
MRMwEQYKCZImiZPyLGQBGRYDY29tMRUwEwYKCZImiZPyLGQBGRYFdm13Y3MxGDAW
<snip>
ukzUuqX7wEhc+QgJWgl41mWZBZ09gfsA9XuXBL0k17IpVHpEgwwrjQz8X68m4I99
dD5Pflf/nLRJvR9jwXl62yk=
-----END CERTIFICATE-----
```

**Private Keys**

Private keys can appear in different formats but are enclosed with clear BEGIN and END markers.

Valid PEM sections begin with one of the following markers.

```
-----BEGIN RSA PRIVATE KEY-----
-----BEGIN PRIVATE KEY-----
```

Encrypted private keys begin with the following marker.

```
-----BEGIN ENCRYPTED PRIVATE KEY-----
```

**Bag Attributes**

Microsoft certificate tools sometimes add Bag Attributes sections to certificate files. VMware Aria OperationsVMware Cloud Foundation Operations safely ignores content outside of BEGIN and END markers, including Bag Attributes sections.

```
Bag Attributes
Microsoft Local Key set: <No Values>
localKeyID: 01 00 00 00
Microsoft CSP Name: Microsoft RSA SChannel Cryptographic Provider
friendlyName: le-WebServer-8dea65d4-c331-40f4-aa0b-205c3c323f62
Key Attributes
X509v3 Key Usage: 10
-----BEGIN PRIVATE KEY-----
```

```
MIICdwIBADANBgkqhkiG9w0BAQEFAASCAmEwggJdAgEAAoGBAKHqyfc+qcQK4yxJ
om3PuB8dYZm34Qlt81GAAnBPYe3B4Q/0ba6PV8GtWG2svIpcl/eflwGHgTU3zJxR
gkKh7I3K5tGESn81ipyKTkPbYebh+aBMqPKrNNUEKlr0M9sa3WSc0o3350tCc1ew
5ZkNYZ4BRUVYWm0HogeGhOthRn2fAgMBAAECgYABhPmGN3FSZKPDG6HJlARvTlBH
KAGVnBGHd0MOmMAbghFBnBKXa8LwD1dgGBng1oOakEXTftkIjdB+uwkU5P4aRrO7
vGujUtRyRCU/4fjLBDuxQL/KpQfruAQaof9uWUwh5W9fEeW3g26fzVL8AFZnbXS0
7Z0AL1H3LNcLd5rpOQJBANnI7vFu06bFxVF+kq6ZOJFMx7x3K4VGxgg+PfFEBEPS
UJ2LuDH5/Rc63BaxFzM/q3B3Jhehvgw61mMyxU7QSSUCQQC+VDuW3XEWJjSiU6KD
gEGpCyJ5SBePbLSukljpGidKkDNlkLgbWVytCVkTAmuoAz33kMWfqIiNcqQbUgVV
UnpzAkB7d0CPO0deSsy8kMdTmKXLKf4qSF0x55epYK/5MZhBYuA1ENrR6mmjW8ke
TDNc6IGm9sVvrFBz2n9kKYpWThrJAkEAk5R69DtW0cbkLy5MqEzOHQauP36gDi1L
WMXPvUfzSYTQ5aM2rrY2/1FtSSkqUwfYh9sw8eDbqVpIV4rc6dDfcwJBALiiDPT0
tz86wySJNeOiUkQm36iXVF8AckPKT9TrbC3Ho7nC8OzL7gEllETa4Zc86Z3wpcGF
BHhEDMHaihyuVgI=
-----END PRIVATE KEY-----
Bag Attributes
localKeyID: 01 00 00 00
1.3.6.1.4.1.311.17.3.92: 00 04 00 00
1.3.6.1.4.1.311.17.3.20: 7F 95 38 07 CB 0C 99 DD 41 23 26 15 8B E8
 D8 4B 0A C8 7D 93
friendlyName: cos-oc-vcops
1.3.6.1.4.1.311.17.3.71: 43 00 4F 00 53 00 2D 00 4F 00 43 00 2D 00
 56 00 43 00 4D 00 35 00 37 00 31 00 2E 00 76 00 6D 00 77 00 61 00
  72 00 65 00 2E 00 63 00 6F 00 6D 00 00 00
1.3.6.1.4.1.311.17.3.87: 00 00 00 00 00 00 00 00 02 00 00 00 20 00
00 00 02 00 00 00 6C 00 64 00 61 00 70 00 3A 00 00 00 7B 00 41 00
45 00 35 00 44 00 44 00 33 00 44 00 30 00 2D 00 36 00 45 00 37 00
30 00 2D 00 34 00 42 00 44 00 42 00 2D 00 39 00 43 00 34 00 31 00
2D 00 31 00 43 00 34 00 41 00 38 00 44 00 43 00 42 00 30 00 38 00
42 00 46 00 7D 00 00 00 70 00 61 00 2D 00 61 00 64 00 63 00 33 00
2E 00 76 00 6D 00 77 00 61 00 72 00 65 00 2E 00 63 00 6F 00 6D 00
5C 00 56 00 4D 00 77 00 61 00 72 00 65 00 20 00 43 00 41 00 00 00
```

```
31 00 32 00 33 00 33 00 30 00 00 00

subject=/CN=cos-oc-vcops.eng.vmware.com

issuer=/DC=com/DC=vmware/CN=VMware CA

-----BEGIN CERTIFICATE-----

MIIFWTCCBEGgAwIBAgIKSJGT5gACAAAwKjANBgkqhkiG9w0BAQUFADBBMRMwEQYK

CZImiZPyLGQBGRYDY29tMRYwFAYKCZImiZPyLGQBGRYGdm13YXJlMRIwEAYDVQQD

EwlWTXdhcmUgQ0EwHhcNMTQwMjA1MTg1OTM2WhcNMTYwMjA1MTg1OTM2WjAmMSQw
```

**Add a Custom Web Certificate to VMware Aria OperationsVMware Cloud Foundation Operations**

Add a Custom Web Certificate
If you did not add your own SSL/TLS certificate when configuring the VMware Aria OperationsVMware Cloud Foundation Operations primary node, you can still add a certificate after VMware Aria OperationsVMware Cloud Foundation Operations is installed.

- Create and configure the primary node.
- Verify that your certificate file meets the requirements for VMware Aria Operations. See the *Getting Started with VMware Aria Operations Guide* or *VMware Aria Operations Installation and Configuration Guide for Linux and Windows*.

1. In a Web browser, navigate to the VMware Cloud Foundation OperationsVMware Aria Operations administration interface at https://*node-FQDN-or-ip-address*/admin.
2. Log in with the admin user name and password.
3. At the upper right, click the SSL certificate icon.
4. In the certificate window, click **Install New Certificate**.
5. Click **Browse for certificate**.
6. Locate the certificate `.pem` file, and click **Open** to load the file in the Certificate Information text box.
7. Click **Install**.

## VMware Aria OperationsVMware Cloud Foundation Operations Passwords

Passwords
For secure VMware Aria OperationsVMware Cloud Foundation Operations operation, you might need to perform maintenance on passwords.

- Passwords are for user access to the product interfaces or to console sessions on cluster nodes.

**Reset the VMware Aria OperationsVMware Cloud Foundation Operations Administrator Password from the Admin UI**

Reset the Administrator Password from the Admin UI
You might need to reset the VMware Aria OperationsVMware Cloud Foundation Operations administrator password as part of securing or maintaining your deployment and if you forget the admin account password.

1. In a Web browser, navigate to the VMware Aria OperationsVMware Cloud Foundation Operations administration interface at https://*<master-node-name> or <master-node-ip-address>*/admin.
2. Log in with the admin user name and password for the master node.
3. In the left pane, click **Administrator Settings**.
4. In the **Change Administrator Password** section, enter the current password, and enter the new password twice to ensure its accuracy.

**NOTE**
You cannot change the administrator user name.

5. Click **Save**.
6. Optionally, to recover a forgotten password, configure the **Password Recovery Settings**.

**Table 310: Password Recovery Settings**

| Password Recovery Settings Options | Description |
|---|---|
| Your E-mail | Email id to which you want to receive the recovery email. |
| SMTP Server | DNS name or IP address of the SMTP server that is used to send the password recovery email. |
| Port | Port used for the communication. By default, 25 is used for a non-secure port and 465 for a secure port. |
| SSL (SMTPS) | Activate to protect the communication using the secure socket layer. |
| STARTTLS Encryption | Activate to switch the insecure communication starting with the TLS handshake. |
| Sender E-mail | The email id from which the password recovery email is sent. |
| User name | User name for the STMP server account, as some servers require authentication. |
| Password | Password for the SMTP server account. |
| Test | To verify the mandatory fields and make an attempt to communicate with the given SMTP server. |

7. Click **Save**. Optionally, click **Reset** to enter the details again.

### Reset the VMware Aria OperationsVMware Cloud Foundation Operations Administrator Password from CLI

Reset the Administrator Password from CLI
You must reset the password if the admin account password is lost.

This procedure requires root account credentials.
- In VMware Aria OperationsVMware Cloud Foundation Operations vApp deployments, when you log in to the console of the virtual application for the first time, you are forced to set a root password.
- The VMware Aria OperationsVMware Cloud Foundation Operations console root password can be different than the admin account password that you set when configuring the VMware Aria OperationsVMware Cloud Foundation Operations primary node.

When the VMware Aria OperationsVMware Cloud Foundation Operations password for the built-in admin account is lost, follow these steps to reset it on vApp clusters.

1. Log in to the master node command-line console as `root`.
2. Enter the following command, and follow the prompts.

```
$VMWARE_PYTHON_BIN $VCOPS_BASE/../vmware-vcopssuite/utilities/sliceConfiguration/bin/
vcopsSetAdminPassword.py --reset
```

### Generate a VMware Aria OperationsVMware Cloud Foundation Operations Passphrase

Generate a Passphrase
When users need to add a node to the VMware Aria OperationsVMware Cloud Foundation Operations cluster, you can generate a temporary passphrase instead of giving them the primary administrator login credentials, which might be a security issue.

Create and configure the primary node.

A temporary passphrase is good for one use only.

1. In a Web browser, navigate to the VMware Aria OperationsVMware Cloud Foundation Operations administration interface at https://*master-node-name-or-ip-address*/admin.
2. Log in with the admin user name and password for the master node.
3. In the list of cluster nodes, select the master node.
4. From the toolbar above the list, click the option to generate a passphrase.
5. Enter a number of hours before the passphrase expires.
6. Click **Generate**.
   A random alphanumeric string appears, which you can send to a user who needs to add a node.

Have the user supply the passphrase when adding a node.

## Give Administrator Access to AD or LDAP Users

You can give users from AD or LDAP administrator access to VMware Aria Operations. To do this, you must log in as the local VMware Aria Operations administrator. Assigning users from AD or LDAP administrative rights will help you distribute the workload among domain users. Once AD or LDAP users get administrative rights, they cannot see the local administrator password, but they can reset it.

**Before You Proceed**

- Know the credentials of the AD or LDAP account.
- Create the group of users who will get adminsitrative acces in AD or LDAP. From VMware Aria Operations, you cannot add or remove any user from the group or view the list of users.

**Procedure**

1. In a Web browser, navigate to the master node administration interface. `https://master-node-name-or-ip-address/admin`.
2. Enter the VMware Cloud Foundation Operations username of the local administrator. `user_name@domain.com`.

   **NOTE**
   AD or LDAP integration in VMware Aria Operations supports all three formats `"DOMAIN\user_name"`, `"CN=username,DC=domain,DC=com"` and `"user_name@domain.com"`. But the Admin UI only supports the email format.

3. Enter the VMware Cloud Foundation Operations local administrator password and click **Log In**.
4. Click **Administrator Settings**.
5. Click the chevron to expand the Active Directory/Open LDAP Integration section.
6. Click the **Activate AD/Open LDAP** button to activate the setting.
7. In the Domain/Subdomain setting, provide the FQDN of the domain name, for example, mydomain.com. Do not provide an IP address. This domain name must be DNS-resolvable.
8. Click the chevron to open Advanced Settings and change the default settings.

**Table 311: Advanced Settings**

| Property | Description |
|---|---|
| Host | The host name is populated to Auto by default. You can select a host from the drop-down list. |
| Port | The port number is populated depending on the **Use SSL/TLS** selection in the domain name. |
| Base DN | The Base Distinguished Name for users is populated based on the domain. Optionally, enter the DN from which to start user searches. For example, cn=Users,dc=myCorp,dc=com. |
| Common Name | The common name is populated to `userPrincipalName` by default. You can change this value from the drop-down menu. |

9. Click the chevron to open Search Criteria and change the default settings.

**Table 312:**

| Property | Description |
|---|---|
| Group Search Criteria | The group search criteria is populated to `(\|(objectclass=group)(objectclass=groupofnames)(objectclass=groupOfUniqueNames))` by default. |
| User Search Criteria | The user search criteria is populated to `(\|(objectclass=user)(objectclass=person)(objectclass=inetOrgPerson)(objectclass=organizationalPerson))` by default. |
| Member Attribute | The member attribute is populated to member by default. |

10. Import the SSL certificate. You can import only one SSL certificate PEM file. The imported SSL certifcate PEM file can contain more than one certificate. For more details, see KB 2046591.
11. Click **Select User Group**. A new dialog box opens.

**Table 313: Select User Group Settings**

| | |
|---|---|
| **Credentials › Username** | Enter the AD/LDAP username. |
| **Credentials › Password** | Enter the AD/LDAP password. |
| Search box | Search for the user group name or distinguished name. |
| Select | Select for the user group name or distinguished name to give administrative access to its users. |

12. Click **TEST** to test if the AD or LDAP connection works.
13. Click **SAVE**.

## OPS-CLI Command-Line Tool

The OPS-CLI tool is a Java application that you can use to manipulate the VMware Aria Operations database. It replaces the `VCOPS-CLI` and `DBCLI` tools.

The product includes the executable file in the tools directory or in <*VCOPS_BASE*>/tools/opscli/.

| Operating System | Filename |
|---|---|
| Linux | `ops-cli.sh` |
| Python | `ops-cli.py` |

All OPS-CLI commands use the `-h` parameter for interactive and localized help.

When you add the `control` command to the `post_install.sh` script, it triggers the redescribe process after an adapter is installed or upgraded.

```
control -h | redescribe --force
```

## Related Command-Line Documentation

In addition to the OPS-CLI, the VMware PowerCLI provides an easy-to-use Windows PowerShell interface for command-line access to administration tasks or for creating executable scripts.

## Supported Operations

The OPS-CLI tool supports the following database operations.

## dashboard Command Operations

You use the `dashboard` command to import, export, share, unshare, delete, reorder, show, hide, and set the default summary for dashboards.

The `dashboard` command uses the following syntax.

```
dashboard -h | import|defsummary|export|share|unshare|delete|reorder|show|hide
[parameters]
```

**Table 314: dashboard Command Options**

| Command Name | Description | Syntax |
|---|---|---|
| dashboard import | Import a dashboard from a file and assign the ownership to a user account. | `dashboard import -h \| user-name\|all\|group:group_name input-file [--force]`<br>`[--share all\|group-name[{,group-name}]] [--retry maxRetryMinutes]`<br>`[--set rank] [--default] [--create]` |
| dashboard export | Export an existing dashboard to a file. | `dashboard export -h \| user-name dashboard-name [output-dir]` |
| dashboard defsummary | Import a dashboard from a file and assign the ownership to a user account. | `dashboard defsummary -h \| input-file\|default`<br>`--adapterKind adapterKind --resourceKind resourceKind` |

*Table continued on next page*

*Continued from previous page*

| Command Name | Description | Syntax |
|---|---|---|
| dashboard share | Share an existing dashboard with one or multiple user groups. | `dashboard share -h \| user-name dahsboard-name all\| group-name[{,group-name}]` |
| dashboard unshare | Stop sharing a dashboard with specified groups. | `dashboard unshare -h \| user-name dashboard-name all\| group-name[{,group-name}]` |
| dashboard delete | Permanently delete a dashboard. | `dashboard delete -h \| user-name\|all\|group:group_name dashboard-name` |
| dashboard reorder | Set the order rank for a dashboard, with an option to make it the default. | `dashboard reorder -h \| user-name\|all\|group:group_name dashboard-name [--set rank] [--default]` |
| dashboard show | Show a dashboard. | `dashboard show -h \| user-name\|all\|group:group_name {,dashbaordname}\|all` |
| dashboard hide | Hide a dashboard. | `dashboard hide -h \| user-name\|all\|group:group_name {,dashboardname}\|all` |

## template Command Operations

You use the `template` command to import, export, share, unshare, delete, and reorder templates.

The `template` command uses the following syntax.

```
template -h | import|export|share|unshare|delete|reorder [parameters]
```

**Table 315: template Command Operations**

| Command Name | Description | Syntax |
|---|---|---|
| template import | Import a template from a file. | `template import -h \| input-file [--force] [--share all\|group-name[{,group-name}]] [--retry maxRetryMinutes] [--set rank] [--create]` |
| template export | Export an existing template to a template file. | `template export -h \| template-name [output-dir]` |
| template share | Share an existing template with one or multiple user groups. | `template share -h \| template-name all\|group-name[{,group-name}]` |
| template unshare | Stop sharing a template with specified groups. | `template unshare -h \| template-name all\|group-name[{,group-name}]` |

*Table continued on next page*

*Continued from previous page*

| Command Name | Description | Syntax |
|---|---|---|
| template delete | Permanently delete a template. | `template delete -h \| template-name` |
| template reorder | Set the order rank for a template. The order rank controls the order of templates created based on shared templates. | `template reorder -h \| template-name [--set rank]` |

## supermetric Command Operations

You use the `supermetric` command to import, export, configure, and delete super metrics.

The `supermetric` command uses the following syntax.
`supermetric -h | import|export|configure|delete [parameters]`

**Table 316: supermetric Command Operations**

| Command Name | Description | Syntax |
|---|---|---|
| supermetric import | Import a super metric from a file and assign the ownership to the specific user account. | `supermetric import -h \| input-file`<br>`[--force] [--policies all\| policy-name[{,policy-name}]]`<br>`[--check (true\|false)] [--retry maxRetryMinutes] [--create]` |
| supermetric export | Export an existing super metric to a template file. | `supermetric export -h \| supermetric-name [output-dir]` |
| supermetric configures | Configure properties of a super metric in one or more super metrics packages. | `supermetric configure -h \| supermetric-name`<br>`--policies all\|policy-name[{,policy-name}]]`<br>`--check (true\|false) --ht (true\|false) --htcriticality level-name`<br>`--dtabove (true\|false) --dtbelow (true\|false)`<br>`--thresholds threshold-def[{,threshold-def}]` |
| supermetric delete | Permanently delete a super metric. | `supermetric delete -h \| supermetric-name` |

## attribute Command Operations

You use the `attribute` command to configure properties of a specific metric in one or more packages. The metric is the object attribute.

The `attribute` command uses the following syntax.

```
attribute configure -h | adapterkind-key:resourcekind-key attribute-key

                         --packages all|package-name[{,package-name}] --check (true|
false)

                         --ht (true|false) --htcriticality level-name

                         --dtabove (true|false) --dtbelow (true|false)

                         --thresholds threshold-def[{,threshold-def}]
```

## reskind Command Operations for Object Types

You use the `reskind` command to configure the default settings in your object type as defined by the ResourceKind model element. The command sets the default attribute or supermetric package, enables or disables dynamic thresholds, and enables or disables early warning smart alerts.

The `reskind` command uses the following syntax.

```
reskind configure -h | adapterkind-key:resourcekind-key

                       --package package-name --smpackage smpackagename

                       --dt (true|false) --smartalert (true|false)
```

## report Command Operations

You use the `report` command to import, export, configure, and delete report definitions.

The `report` command uses the following syntax.

```
report -h | import|export|delete [parameters]
```

**Table 317: report Command Options**

| Command Name | Description | Syntax |
|---|---|---|
| report import | Import a report definition from a file. | `report import -h \| input-file [--force]` |
| report export | Export one or more report definitions to a file. | `report export -h \| all\|report-name[{,report-name}] [output-dir]` |
| report delete | Permanently delete one or more report definitions. | `report delete -h \| all\|report-name[{,report-name}]` |

## view Command Operations

You use the `view` command to import, export, or delete view definitions.

The `view` command uses the following syntax.

```
view -h | import|export|delete [parameters]
```

**Table 318: view Command Operations**

| Command Name | Description | Syntax |
|---|---|---|
| view import | Import a view definition from a file. | `view import -h \| input-file [--force]` |
| view export | Export one or more view definitions to a file. | `view export -h \| all\|view-name[{,view-name}] [output-dir]` |
| view delete | Permanently delete one or more view definitions. | `view delete -h \| all\|view-name[{,view-name}]` |

## file Command Operations

You use the `file` command to import, export, list, or delete database files. The command operates on metric, text widget, and topology widget files.

The `file` command uses the following syntax.

```
file -h | import|export|delete|list [parameters]
```

**Table 319: file Command Operations**

| Command Name | Description | Syntax |
|---|---|---|
| file import | Import a metric or widget from a file. | `file import -h \| reskndmetric\| textwidget\|topowidget`<br><br>`input-file [--title title] [--force]` |
| file export | Export one or more metrics or text widgets, or export the topology widget to a file. | `file export -h \| reskndmetric\| textwidget\|topowidget`<br><br>`all\|title[{,title}] [output-dir]` |
| file delete | Permanently delete a metric or a widget. | `file delete -h \| reskndmetric\| textwidget\|topowidget`<br><br>`all\|title[{,title}]` |
| file list | List all metric or a widget files. | `file list -h \| reskndmetric\| textwidget\|topowidget` |

# VMware Aria Operations User Guide (8.18)

Use VMware Aria OperationsVMware Cloud Foundation Operations to automate and manage your IT with full stack visibility from the physical, virtual and cloud infrastructure to the applications they support.

VMware Aria Operations delivers intelligent operations management with application-to-storage visibility across physical, virtual, and cloud infrastructures. Using policy-based automation, operations teams automate key processes and improve the IT efficiency.

Using data collected from system resources (objects), VMware Aria Operations identifies issues in any monitored system component, often before the customer notices a problem. VMware Aria Operations also frequently suggests corrective actions you can take to fix the problem right away. For more challenging problems, VMware Aria Operations offers rich analytical tools that allow you to review and manipulate object data to reveal hidden issues, investigate complex technical problems, identify trends, or drill down to gauge the health of a single object.

Using these capabilities of VMware Aria OperationsVMware Cloud Foundation Operations, as a system administrator, you become aware of a problem with an object in your environment when VMware Aria OperationsVMware Cloud Foundation Operations generates an alert, or when a user contacts you. To help ensure optimal performance, this information describes how you use VMware Aria OperationsVMware Cloud Foundation Operations to monitor, troubleshoot, and take action to address problems. It also provides information on how to assess whether problems due to over demand or lack of capacity require a system change or upgrade.

## Intended Audience

This information is intended for VMware Aria OperationsVMware Cloud Foundation Operations administrators, virtual infrastructure administrators, and operations engineers who track and maintain object performance in your managed environment.

## Managing your Environment Configurations

You can use VMware Cloud Foundation Operations Configuration Drifts feature to monitor and view vCenter configuration settings that have drifted from assigned templates without needing to track every change manually.

Configuration drift shows the changes in product configuration over time and allows you to compare the changes to the assigned template values. It helps to prevent misconfiguration from going unnoticed, reduces the risk of security breaches, and keeps the environment running smoothly.

Configuration templates allow administrators to define and review specific configuration settings for vCenter instances. The desired state can be defined as a configuration template or JavaScript Object Notation (JSON) based file that contains settings for vCenter instances, such as network configuration, storage, security, advanced settings, and performance.

> **NOTE**
> This feature applies to only vCenter 8.0.3 version or later.

## Why use Configuration Drift

Configuration Drift offers several benefits, including:

1. Consistency: By defining a template, administrators can ensure that a specific group of vCenter instances or all of them are configured consistently. This helps reduce errors and improves the overall reliability of the virtual infrastructure.
2. Compliance: Configuration Drift helps ensure that vCenter instances are configured in compliance with organizational policies and regulations.
3. Scalability: As infrastructure environment scales up, managing configuration drifts becomes crucial for maintaining consistency and control across global on-premise vCenter instances.

## Where to begin

This overview provides a general guide for global administrator or local administrator on where to begin.

**Day in the life of administrator**

A global or a local admin can perform these tasks to define global standard configurations across organization and monitor the configuration settings.

- Creating a configuration template.
- Detecting drifts as compared to assigned templates.
- Viewing and monitoring all vCenter instances and their configuration specifications in one place.

## Managing Config Templates with Version Control Systems

Integrating VMware Cloud Foundation Operations with Version Control System (VCS), allows you to control versioning, auditing, and the automatic management of all your configuration templates. This integration supports the Git based Version Control Systems like GitLab and GitHub.

**Integrating with Git**

Integrating with Version Control Systems like GitLab or GitHub repositories enable you to manage the configuration templates under source control. This functionality also facilitates accurate configuration updates, supports change control with reviews, and integrates CI/CD for automated testing and validation. For more information on Git, please see About Git.

**Before You Begin**

- Verify that your VMware Aria Life Cycle fleet management can access the GitLab or GitHub server.
- Generate access tokens to authenticate and authorize a secure connection to the repositories.

  **GitLab**: For GitLab, generate an access token for your GitLab project and copy it to your clipboard for use during the configuration process, see Creating a personal access token.

  **NOTE**
  To generate the access token for GitLab, you must have the Maintainer or Owner role. When you generate an access token for GitLab, under **Access Tokens**, enter a name for the token, set the expiration date, and select the appropriate scopes:

  - `read_repository`: For read-only access to the repository.
  - `write_repository`: For write access to the repository.
  - `read_api`: For access to the container registry.
  - `api`: For complete read and write access to the scoped API project that includes the package registry.

  **GitHub**: For GitHub, you must use a GitHub repository and generate a GitHub token, see Creating a personal access token for the command line.

**Configuring Source Control in VMware Cloud Foundation Operations**

You can add GitLab or GitHub as the Source Control in VMware Cloud Foundation Operations to sync config templates from VMware Cloud Foundation Operations to the Git remote repository.

Integrate Git in VMware Cloud Foundation Operations:

1. From the left menu, click **Administration › Control Panel** and then click **Source Control**.

2. For **Source Control Type**, select **GitLab** or **GitHub**.
3. Enter the **Source Control Server** address where the configuration templates will be stored.
4. Enter a valid **Access Key** with the following permissions:
   - GitLab: Create a personal access token with all scopes selected.
   - GitHub: Create an access token with the following scopes selected: `api`, `read_api`, `read_repository`, and `write_repository`.
5. Specify the **Branch** for storing or reading content.
6. Provide the **Repository** path for storing or retrieving contents.

   A repository is a directory initialized by Git to store files along with their version history. For more information, see Getting a Git Repository.

7. By default, all config templates directly sync to the remote repository without a review process. Click the **Code Review** check box to enable a review process before changes are committed.

   If you select **Code Review**, a merge request gets created for each config template. The content syncs to the repository once the merge request is approved.

8. Click **Test Connection** to verify if the integration is successful. If the **Connection Result** is successful, click **Save** to complete the integration.

Once the Git integration is complete, you can manage the config templates from GitLab or GitHub.

**Template Organization and Repository Structure**: Config templates are automatically created and organized in a predefined folder structure: `Operations-ConfigTemplate/config-templates/vCenter`. Each template must stored in a uniquely named folder within the repository.

**Template Name and Storage**: Each config template must have a unique name in the repository to avoid duplication.

**Template Content and File Structure**: For each config template, three files are automatically created in the repository:

1. `<templatename>-desiredState.json`: Contains the actual desired state of the template.
2. `<templatename>-metaData.json`: Contains metadata information like template name, version, selected configurations, and other relevant details.
3. `<templatename>-manifest.json`: Used to manage and maintain the template's structure and metadata.

To successfully integrate config templates with GitLab or GitHub, you must configure your GIT profile in VMware Cloud Foundation Operations and then sync your config templates. This integration ensures that the templates are automatically created, organized, and managed effectively in the repository, with proper access control at the repository level.

### Syncing Config Templates with Git Integration

After you configure Git successfully in VMware Cloud Foundation Operations, sync the configuration templates with the Git repository. This topic outlines the different workflows you can follow to manage the config templates and integrate them with a remote Git repository.

#### Create, Edit, or Delete a Template from VMware Cloud Foundation Operations

**Review Process Enabled**

When you create, edit, or delete a template in VMware Cloud Foundation Operations, a **Merge Request** is automatically created and visible in the Config Template list view.

You must review and approve the merge request in GitLab

After the merge request is approved, the template status changes to **OUT_OF_SYNC**.

Sync the template using the **Pull from GitLab/GitHub** option in VMware Cloud Foundation Operations.

The template status changes to **IN_SYNC** after the sync is complete.

> **NOTE**
> The edited or deleted content is reflected only after the merge request is approved and the content is synced.

**Review Process Disabled**

When you create or edit a template, the template is directly created or edited in the remote Git repository without a merge request. If you delete a template from VMware Cloud Foundation Operations, the template is marked **OUT_OF_SYNC** and is directly deleted from the remote repository during the next sync.

## Sync Existing Templates in VMware Cloud Foundation Operations

**Review Process Enabled**

When Git integration is enabled, **Merge Requests** are automatically created for syncing all existing templates from VMware Cloud Foundation Operations to the remote repository. After the merge request is approved, the template status changes to **OUT_OF_SYNC**. Use the **Pull from GitLab/GitHub** option to sync content from the remote repository. The status of the templates changes to **IN_SYNC** once the sync is successful.

**Review Process Disabled**

The status of all existing templates reflects as **OUT_OF_SYNC**: No merge request is raised. You must use the **Pull from GitLab/GitHub** option to sync content from the remote repository. The status changes to **IN_SYNC** after the sync is complete.

## Sync Remote Templates to VMware Cloud Foundation Operations

If you create, edit, or delete templates in the Git repository, the status of the corresponding templates reflects as **OUT_OF_SYNC** in VMware Cloud Foundation Operations.

If you create a new template in the remote repository, an empty placeholder template with the same name is created in VMware Cloud Foundation Operations. You must use the **Pull from GitLab/GitHub** option to sync content from the remote repository. The status changes to **IN_SYNC** after the sync is complete.

If you edit a template in the remote repository, you must **Pull from GitLab/GitHub**, and once the syc is successful, the updated template content gets reflected in VMware Cloud Foundation Operations and the status changes to **IN_SYNC**.

If you delete a template in the remote repository, the template is removed from VMware Cloud Foundation Operations.

If there are existing templates in the remote Git repository, when you configure Git in VMware Cloud Foundation Operations, templates without any content get created in VMware Cloud Foundation Operations as placeholders and reflect the **OUT_OF_SYNC** status. You must sync the content to populate the templates. The status changes to **IN_SYNC** after the sync is complete.

## Additional Information

**Merge Request Discarded**

A merge request can be discarded if the changes proposed in the merge request are no longer valid or required. If a merge request is discarded or closed by the user, a warning message is displayed informing the user that the changes made to the corresponding template in VMware Cloud Foundation Operations are not merged and the current template may not reflect the desired state in the remote Git repository.

You can edit the template to make the necessary updates and address the issues that caused the merge request to be discarded. After you make the update a new merge request is raised. This ensures that the latest template changes are submitted for review.

**Drift Detection**

When drift detection does not run for a specific template, the drift checks are skipped if the template does not exist in the remote Git repository. This may occur in scenarios where the template has not synced with the remote repository yet.

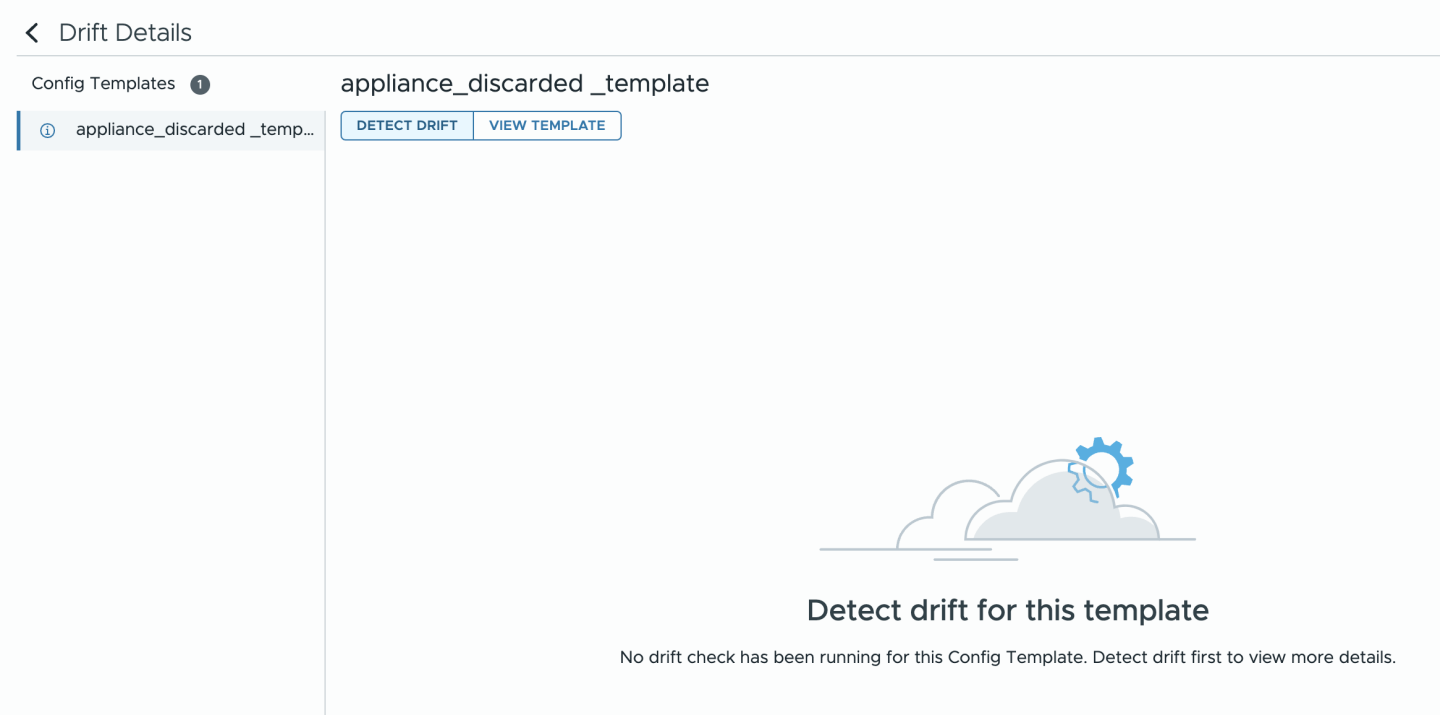


**Updating Git Configurations**

**Editing or Updating Git Branch or Git Repositories**

You can update the Git branch or repository or switch to a new branch or repository. Once done, the templates from the old repository display the **OUT_OF_SYNC** status. You must only sync the existing templates in VMware Cloud Foundation Operations to the new branch or repository.

If you delete a Git configuration, the templates stored locally in VMware Cloud Foundation Operations remain unchanged.

If you edit, update or delete configurations, do not sync or update the templates that were **OUT_OF_SYNC**.

> **NOTE**
> If you want to update your source control type, you must first delete the existing source control type and then configure the new one.

## Viewing and Detecting a Configuration Drift

The Drifts feature allows you to view and detect configuration drifts from assigned vCenter Configuration Templates.

To view vCenter configuration setting values and detect configuration drifts, you must:
- Ensure that you have the **Administrator** role.

- Verify that you have **Configuration Drifts** privileges to access configuration templates. For more information, see the Managing Users and Access Control in VMware Aria Operations topic in the *Configuring VMware Aria Operations* guide.
  - **Manage Privilege** : for creating templates, assigning them to policy, and running drift check.
  - **View Privilege**: for viewing template content and viewing existing drift.
- Verify that you have vCenter configured and registered with the cloud proxy.
- Verify that you have one active policy assigned to the configuration template.

On the Drifts page, the Configuration Drifts table lists all the vCenter instances.

| Option | Description |
|---|---|
| vCenter | Displays the vCenter instances. |
| Drift Status | Displays how many configuration drifts were detected. |
| Templates | Displays the number of configuration templates associated with the vCenter instance. |
| Last Drift Check | Displays the timestamp of the most recent drift detection. |

Click **Manage Columns** to choose which columns to display in the table.
To detect a configuration drift:

1. From the left menu, click **Management** > **Configuration Management** > **Drifts**.
2. To detect drift for a vCenter instance, select the instance and click **Detect Drift**.

   This detects drifts across all the configuration templates associated with the vCenter instance.

The drift status is displayed in the Drifts table.

**Scheduling Drift Detection**

You can now automate the process of drift detection and schedule drift detection in VMware Cloud Foundation Operations.

Scheduling Drift Detection enables you to create a drift detection job to run automatic drift checks on your vCenter instances at an interval of your choice.
To schedule drift checks:

1. From the left menu, click **Management** > **Configuration Management** > **Schedule Drift Detection**.
2. Enter **Configuration Drift Check Information** and click **Next**.

   You can schedule drifts only for vCenter object types.

3. **Select an object**, **Select Scope** from the list of vCenter instances. If you select a VCF folder as the scope, all VCF instances belonging to the folder get automatically added to the scope.
4. Click **Preview Scope** to confirm the vCenter instances that the drift check will run on, and then click **Next**.
5. You can filter the criteria and add multiple criteria for the vCenter object type and then click **Next**.
6. Set the **Schedule** and click **Create**.
   A new job is created to run scheduled drift detection in the automation central page. For more information see the 'Automation Central' topic in the *Configuring VMware Cloud Foundation Operations guide.*

   You can download the drift report to view the historical data. By default the report shows data for the last seven days. You can update the duration to up to 30 days from the Global Settings page. For more information, see the 'List of Global Settings' topic in the *Configuring VMware Cloud Foundation Operations guide.*

### Viewing Drifts for the vCenter Instance

You can detect configuration drifts for each configuration template you have created for a specific vCenter instance.

Ensure that you have created at least one configuration template for your vCenter Instance.

To view your configuration drift details:

1. On the Drifts page, click **vCenter** instance.
2. On the **Drifts Details** page, you can view the number of configuration templates assigned with the vCenter instance. Select a configuration template to view the drift.
3. Click **Detect Drift** to detect configuration drift. The right pane displays the drifts between the **Current vCenter Value** and the **Config Template Value**.

   **NOTE**
   If no drift is detected for the selected template, the message **No drift found for this Config Template since last configuration drift check** status is displayed.

4. Click **View Template** to view desired configuration settings for the selected vCenter.

### Viewing Drifts for Clusters

You can view the drift status of your VCF clusters in VMware Cloud Foundation Operations. The drift status is visible only for clusters that have vSphere Configuration Profile enabled.

- Verify that you have vCenter version 9.0 or later.
- Verify that you have registered your vCenter account with the cloud proxy.

  **NOTE**
  Select the Cloud Proxy deployed on the given vCenter. If the outbound internet access for the cloud proxy must be restricted, ensure that the minimum Cloud Proxy prerequisites are met.

  For more information, see the Configuring Cloud Proxies in VMware Cloud Foundation Operations topic in the *Configuring VMware Aria Operations* guide.

- Ensure that you configured at least one VCF account in VMware Cloud Foundation Operations. For more information, see the "Configuring VMware Cloud Foundation Cloud Account" topic in the VMware Cloud Foundation Operations Configuration Guide.
- Clusters are managed via configuration profiles in individual vCenter instances. Enable vSphere Configuration Profiles for clusters in vSphere.

  **NOTE**
  There may be a delay of up to 8 hours for updates made on the cluster in vCenter instances to be reflected on the Configuration Drifts page

To view the cluster drift status:

1. From the left menu, click **Fleet Management> Configuration Drift**.

   The Cluster Drift Status tile displays the drift status of all the clusters in the VCF account. You can also view the cluster drift status of individual VCF instances or individual domains within the instances.

2. Click a VCF instance or domain and then click **View All Cluster Drifts** to view the cluster information and also see if a cluster is enabled or disabled for drift detection.

**NOTE**
You can click on a cluster to view the cluster profile in vSphere.

### Using the Configuration Drift Dashboard

The Configuration Drift dashboard allows you to view and monitor all vCenter instances and their configuration specifications in one place.

You can view and compare configuration drifts, identify changed settings, and determine the root cause of problems. Additionally, you can view the configuration templates assigned to your vCenter instances.

A configuration drift notifies you of any deviation from assigned template values. **vCenter by Drift Status** pie-chart represents the state of the vCenter at a given point in time. The total number of vCenter drift statuses are:

- **Compliant**: This state shows vCenter with no drifts against the template it has been associated with.
- **Non-Compliant**: This state shows whether the vCenter configurations deviate from the desired standard configurations.
- **Unavailable**: This state occurs when vCenter is unavailable, the drift computation is in progress, the drift has not been computed against that specific resource, or an internal error has occurred.
- **Not-Supported**: This state shows the list of vCenter instances that are either below version 8.0.3 or are not registered with the cloud proxy.

To see the number drifts for each vCenter, click **View Drifts**.

A template enables you to apply the standard configurations that you have defined across clusters, vCenter instances, SDDCs, etc. **vCenter by Templates** displays the number of configuration templates associated with vCenter instances.

- **With Templates**: This state shows the number of vCenter instances that have at least one template associated with them.
- **No templates**: This state shows the number of vCenter instances that have no template associated with them.
- **Not-Supported**: This state displays the list of vCenter instances that are not registered with the cloud proxy or are below version 8.0.3.

To view the templates created for each vCenter, click **View Templates**.

## Troubleshooting Configuration Drift

Use this troubleshooting topic to remediate commonly encountered problems with configuration drift

**Configuration Drifts Internal Server Error**

**Issue 1:** An internal server error occurred while displaying drift details or retrieving Configuration Settings from the selected vCenter instance.

Possible Causes:

- vCenter Server is not reachable from Cloud Proxy.
- The Configuration Management Adapter is in an ERROR state. (See Appendix on how to check Configuration Management Adapter status.)
- Configuration Module service is not running. (See Appendix on how to check if Configuration Module is running.)

Solution:

- Fix the connectivity of vCenter instance from Cloud Proxy and ensure vCenter services are up. SSH into Cloud Proxy and verify the connectivity and verify the connection inside the cloud proxy by pinging `<vc_ip>`.
- Fix the ERROR state in the Configuration Management Adapter. For more information, see Cloud Proxy Troubleshooting.

- SSH into Cloud Proxy and restart Configuration Module docker service: `docker restart config-modules-docker`.

**Appendix:**

- Check if the Configuration Management Adapter is “Healthy”.
- Check if vCenter instance is reachable and in “Collecting” state.
- Check if the Configuration Module service is running.

1. Log into the VMware Cloud Foundation Operations cloud proxy node as root via SSH
2. Run the following command to see the status of the Configuration Modules service:docker ps

```
root@localhost [ ~ ]# docker ps
CONTAINER ID   IMAGE                                                                                COMMAND
          CREATED        STATUS        PORTS
                                  NAMES
e2a182ba6aeb   cmd-docker-local.artifactory.eng.vmware.com/config-modules-docker:1.1.8          "./scripts/sta
rt_api…"   12 hours ago   Up 12 hours   443/tcp
                                  config-modules-docker

198ec0671e97   ucp-docker-local.artifactory.eng.vmware.com/salt-master-bootstrap:8.18.0.20187   "/bin/sh -c '/
tmp/sa…"   12 hours ago   Up 4 hours    0.0.0.0:4505-4506->4505-4506/tcp, :::4505-4506->4505-4506/tcp, 0.0.0.0
:8553->8553/tcp, :::8553->8553/tcp   ucp-controlplane-saltmaster
root@localhost [ ~ ]#
```

**Cluster Configuration**

**Issue 1:** ESXi Clusters arenot showing up on the Drifts Page.

Solution: Enable vSphere Configuration Profile on the cluster. For more information, see Configuration Management using vSphere Configuration Profile.

**Issue 2:** Inconsistency in the vSphere Configuration Profile drift status on the vCenter Console and Admin Console.

Solution: The vSphere Configuration Profile status is polled every 8 hours. In case of inconsistency, wait few hours for the next collection cycle to complete.

**Issue 3:** No cluster-related data displayed on the Objects tab, Drifts tab, or the Overview page on any VCF, SDDC, or Domain level.

Solution: Confirm if there are any clusters in your inventory. If there are clusters in your inventory, ensure you have registered your SDDC stack. It usually takes 5-6 minutes for the inventory to load and push the data for resources including vCenter instances or clusters. Check the Inventory tab to confirm if that Inventory is loaded properly.

**Configuration Version Control**

**Issue 1:** Template Sync Failure from Local to Remote or Remote to Local.

When a template fails to sync to the remote repository or from the remote repository, error information will be displayed at the template level in VMware Cloud Foundation Operations.

Possible Cause 1: vRLCM Not Reachable: VMware Cloud Foundation Operations cannot communicate with the vRealize Lifecycle Manager (vRLCM).

Solution:

1. Check the status of all services running on the vRLCM machine and ensure they are active.
2. SSH to vRLCM machine.
3. Check the status of vRLCM service is `enabled`.

4. If the vRLCM service is not running. Start the service `service vrlcm-server start`.
5. Ensure that the service is running and the the service status is **enabled** .

Possible Cause 2: Remote Repository Not Reachable: vRealize Lifecycle Manager (vRLCM) cannot establish a connection to the configured remote Git repository.

Solution:

- Ensure that the vRLCM machine has network connectivity to the remote repository server (GitLab or GitHub).

1. Ping the Remote Repository Server: `ping <repository_server_address>Replace <repository_server_address>`
2. Replace `<repository_server_address>` with the GitLab or GitHub server address (github.com or gitlab.com).
3. Ensure that the ping requests are reaching the server. If not, there may be a network issue blocking the connection.

- Verify that any firewalls or network policies are not blocking access to the repository server, trace route the Server: `traceroute <repository_server_address>`. This helps identify where the connection is failing (e.g., firewall or proxy).

Possible Cause 3: Invalid Remote Repository Credentials: The provided access token for the remote repository is invalid or lacks the necessary permissions.

Solution:

1. Confirm that the access token provided in the source control configuration is correct.
2. Ensure the token has the following permissions: Read repository contents, Push changes, Create and manage merge requests and Trigger notifications.

**Issue 2:** Invalid Template Format

During syncing the template content from the Remote repository the Sync failed because the template format is not supported or is missing required fields, causing the sync to fail.

Possible Cause 1: Corrupted Template Content: The template content has syntax errors or is in an unsupported format.

Solution: Validate the template content against the required schema and correct any formatting or syntax errors.

Possible Cause 2: Unsupported Template Version: The template was created in an older version that is no longer supported.

Solution: Update the template to the latest supported version in the GIT repository directly.

1. Go to repository and then go to the template folder.
2. Open `<<template>>-metadata.json`
3. Update the version with the vCenter instance version available in VMware Cloud Foundation Operations.

**Issue 3:** Access Token Expired: Sync operations fail due to an expired access token.

Possible Cause: The access token used for the remote repository has expired.

Solution:

1. Navigate to **Administration › Control Panel › Source Control**.
2. Verify **Test Connection** is successful with provided configuration.
3. You can also verify in the GItlab/Github Personal Access token section to check expiration status.
4. If the test fails, generate a new access token with the required permissions.

5. Update the source control configuration in VMware Cloud Foundation Operations with the new token.

**Error with Git Repository**

**Issue 1:** Repository Not Reachable

Troubleshooting steps:

1. Navigate to **Administration › Control Panel › Source Control**.
2. Click **Test Connection** to verify access to repository.

Solution: Ensure that the vRLCM machine has network connectivity to the repository server and check the firewall and DNS settings to ensure there are no restrictions blocking access. Once done, **Test Connection** to confirm that the connection to the repository is successful.

**Issue 2:** Access Token Expired

Troubleshooting steps:

1. Navigate to **Administration › Control Panel › Source Control**.
2. Click **Test Connection**.
3. Review error message indicating token expiry.

Solution: Generate a new access token from the Git repository with required permissions and update the token in theVMware Cloud Foundation Operations . Once done, **Test Connection** to confirm that the updated token is working. You can also perform a sync operation to confirm the token validity.

**Issue 3:** Invalid Access Token Permissions

Troubleshooting steps:

1. Navigate to **Administration › Control Panel › Source Control**.
2. Review the error message indicating missing permissions.

Solution: Update the access token to include required permissions, such as Read repository, Push changes, and Manage merge requests and then update the token in VMware Cloud Foundation Operations. Once done, perform any configuration management operations like creating or editing a template to ensure that the sync operation completes without any permission errors.

**Issue 4:** Branch Not Available

Troubleshooting steps:

1. Navigate to **Administration › Control Panel › Source Control**.
2. Verify the branch name and check if the branch exists in the remote repository.

Solution: Create the missing branch in the remote repository and update the branch name of the existing branch in VMware Cloud Foundation Operations. Once done, **Test Connection** to confirm that the branch is accessible. You can also perform a template sync to confirm that the branch is valid.

**Issue 5:** Merge Request Fails to Create

Troubleshooting steps:

1. Navigate to **Administration › Control Panel › Source Control**.
2. Click **Test Connection**.
3. Review error message when attempting to create a merge request.

Solution: Ensure that the configured branch allows merge requests and also check for repository restrictions like branch protection rules and resolve them. Verify token permissions for managing merge requests and then try to create or edit a template with review process enabled to validate that the merge request is created successfully.

**Issue 6:** Remote Repository Authentication Failed

Troubleshooting steps:

1. Navigate to **Administration › Control Panel › Source Control**.
2. Click **Test Connection**.
3. Review error message indicating authentication failure.

Solution: Ensure the provided access token is valid and update the access token in VMware Cloud Foundation Operations. Verify that the token belongs to a user with access to the repository. Once done, **Test Connection** to validate that the authentication is successful.

**Alarm and Notification**

**Issue 1:** Alert notification is not generated after the vCenter drifts against a template either with status as "Drifted" or "Error".

Solution: Please wait for the next collection cycle.

**Issue 2:** Notification is received for vCenter being drifted against a template.

Solution:

1. Log on to vCenter and revert the drifted configurations to its original value to make the vCenter complaint again.
2. Under **Infrastructure Configuration,** edit the template to revert the drifted configurations to its original value.

**Issue 3:** Error-ed drifted state.

Solution: Check if there is some connection error between VMware Cloud Foundation Operations and vCenter, for example, confirm whether the vCenter instance is collecting data or not. If not, start the collection.

**Configuration Drift Reports**

**Issue 1:** When the drift report PDF gets generated with "Unable to generate configuration drifts report currently, please try again later."

Solution: If any internal issue occurs then a PDF gets generated with this message. Please try again after some time and please check the exception logs in this class to fix the issue.

`Class: com.vmware.vrops.config.util.PdfGenerationUtil`and look for this exception message `"Error caught while creating PDF error"`.

**Issue 2:** When drift report PDF gets generated with "No drift details available for the selected resource/resources."

Solution: Confirm if **Detect Drift** was triggered for selected resources. Trigger drifts to the selected resources and download the report again.

## What is Diagnostics for VMware Cloud Foundation

Diagnostics for VMware Cloud Foundation is a centralized platform that monitors the overall operational status of the VMware Cloud Foundation software stack.

It is a self-service platform that helps you analyze and troubleshoot the components of VMware Cloud Foundation, including vCenter, ESXi, vSAN, capabilities such as vSphere vMotion, snapshots, VM provisioning, and other issues including security advisories and certificates. As an Infrastructure admin, you can monitor the operational state of your

environment using diagnostics findings and custom dashboards. The built-in dashboards are an extension to native VMware Cloud Foundation Operations dashboards. Diagnostics validates if your environment is up-to-date with the important VMware Security Advisories.

With Diagnostics, you can address issues or vulnerabilities related to certificates such as expired SSL certificates.

Diagnostics also provides relevant information in self-help flows for vCenter capabilities such as vSphere vMotion to help you diagnose migration issues.

**Key Benefits**:

1. Ensures platform availability by proactively identifying and diagnosing operational issues.
2. Preserves the security posture of your environment.
3. Provides built-in known issue detection with remediation guidance and links to supporting Knowledge Base articles.
4. Self-service improves the time to understand the cause of an issue and determine the next steps for your VMware software environment.
5. Quick identification of the cause and remediation options helps your business run with less disruption.

Diagnostics monitors the operational state of the following components, which also reports findings indicating the occurrence of a known issue and provides recommendations:

- vCenter: vCenter Operational State: Ping reachability
- ESXi Operational State: Connectivity from vCenter
- VMware vSAN Operational state: Disk group status, Physical disks status
- General issues: Certificate expiry

Diagnostics includes the following set of capabilities that help understand the operational state of the VMware Cloud Foundation software stack:

- Workload Provisioning (VMs): VM Provisioning requests and failures, provisioning findings and recommendations, and general troubleshooting.
- vSphere vMotion : Successful and failed vMotion, findings and recommendations, and general troubleshooting.
- Snapshots: Snapshot failures, findings and recommendations, and general troubleshooting.

The Diagnostics for VMware Cloud Foundation uses interactive cards to display data. Click **View Details** or **View Dashboard** for more information. See the topic Working with VMware Cloud Foundation Diagnostics.

Diagnostics cards with content appear if your environment has VMware Cloud Foundation Operations integrations configured. To configure the VMware Cloud Foundation Operations integrations , follow the instructions in Setting up VMware Cloud Foundation Diagnostics.

## How Diagnostics for VMware Cloud Foundation works

Diagnostics for VMware Cloud Foundation consolidates signatures from VMware Skyline Advisor and VMware Skyline Health Diagnostics, and integrates VMware Cloud Foundation Operations for logs to provide a single pane for monitoring and troubleshooting. The issues that Diagnostics detects are called findings. Diagnostics scans system properties and product logs, and displays findings that you can act upon. You review the findings and decide on the next steps appropriate for your VMware software environment. Findings are different from reports on the operational state of a system, such as connectivity, services status or interface issues. Property-based findings inform you about issues that might affect your environment. Log-based findings inform you if an issue has already affected your system. Diagnostics 5.2 works with more than 300 property and log-based rules. You can see a list of all signatures here.

When you experience an issue in your environment, you can initiate a log scan within Diagnostics which uses existing signatures to detect issues. When a signature matches the information in the log files, findings are displayed. The finding contains information about the matching signature and remediation steps or a Knowledge Base article to help resolve the issues.

Proactive findings are based on rules that inspect system properties using APIs. These rules are run automatically every four hours. To detect issues that have already occurred in your environment, you can initiate a log scan by clicking **Refresh Findings**. To run a log scan, you must have VMware Cloud Foundation Operations for logs installed and integrated in your environment. See, Setting up Diagnostics for VMware Cloud Foundation.

### Architecture Diagram and Data Flow of Diagnostics for VMware Cloud Foundation

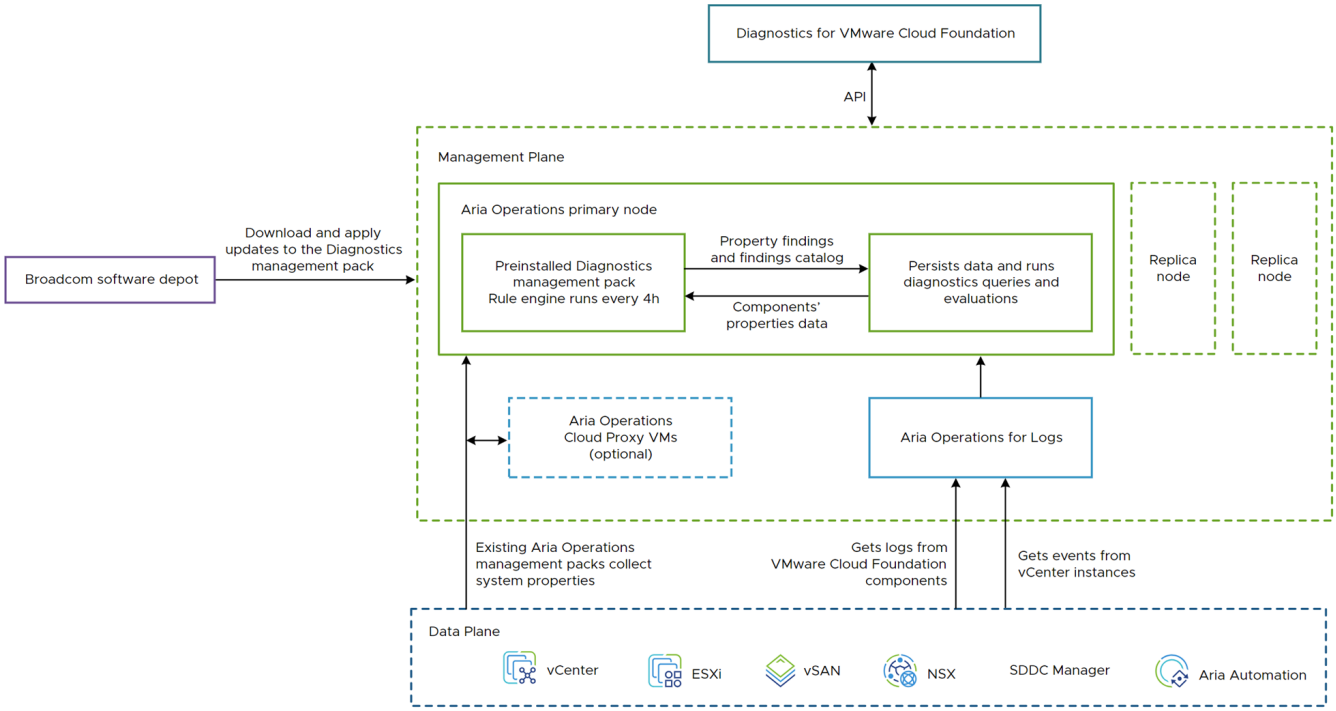


See VMware Cloud Foundation Operations release notes to get updates for the management packs.

## How you discover data in the Diagnostics dashboard

Your self-help flow on Diagnostics can start either from one of the cards or from Overall Findings. The starting point for troubleshooting depends on how the issue is reported or identified. You might see triggers for investigation and corrective action in the Diagnostics dashboard or receive a report from an external source, such as an issue reported from an end-user who is not an infrastructure administrator.

Look at the Diagnostics Card Self-help Flow and Diagnostics Findings Self-help Flow flowcharts for more details.

### Diagnostics Card Self-help Flow

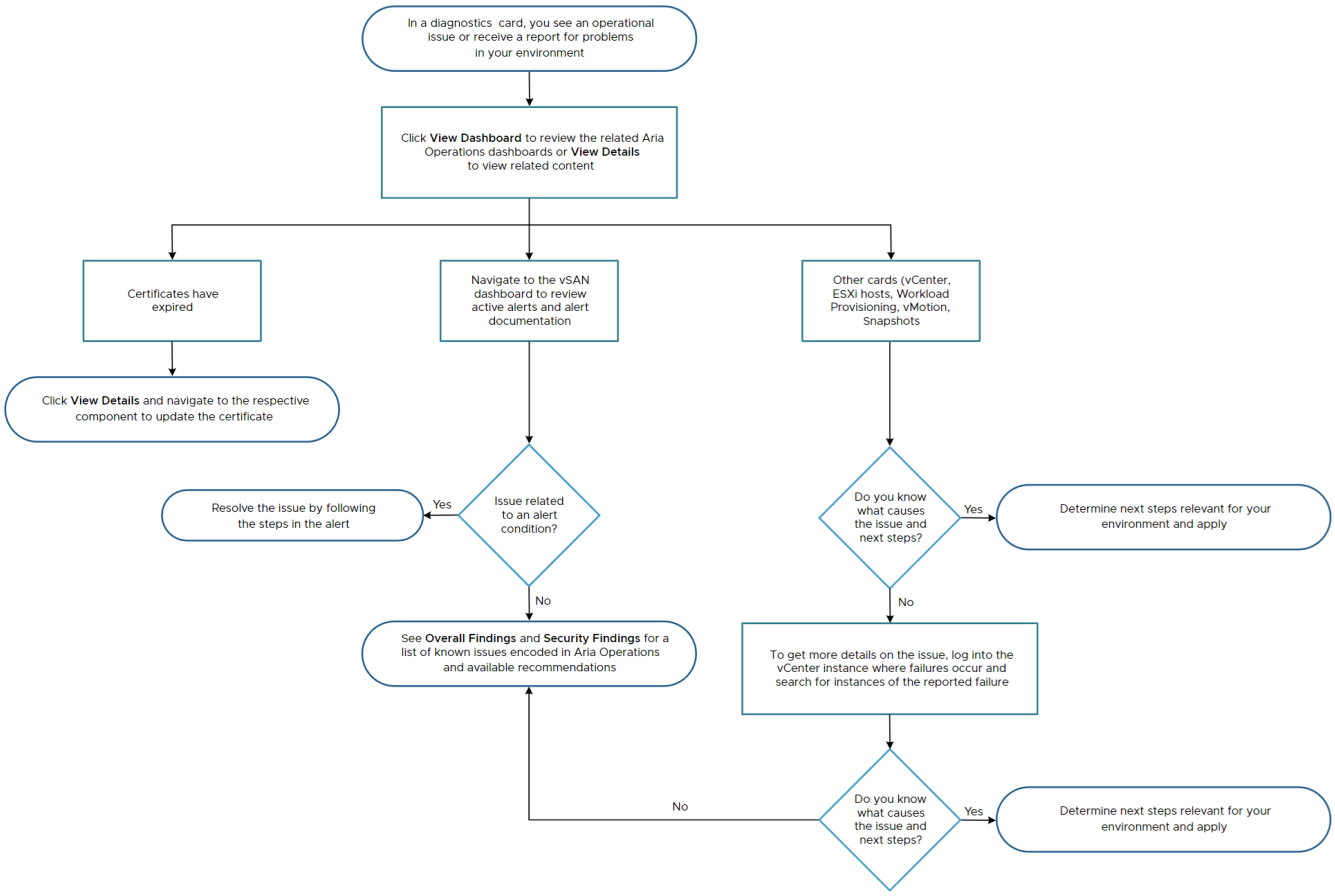
In a diagnostics card, you see an operational issue or receive a report for problems in your environment
Click View Dashboard to review the related Aria Operations dashboards or View Details to view related content
Certificates have expired
Navigate to the vSAN dashboard to review active alerts and alert documentation
Other cards (vCenter, ESXi hosts, Workload Provisioning, vMotion, Snapshots
Click View Details and navigate to the respective component to update the certificate
Issue related to an alert condition?
Yes
Resolve the issue by following the steps in the alert
No
See Overall Findings and Security Findings for a list of known issues encoded in Aria Operations and available recommendations
Do you know what causes the issue and next steps?
Yes
Determine next steps relevant for your environment and apply
No
To get more details on the issue, log into the vCenter instance where failures occur and search for instances of the reported failure
Do you know what causes the issue and next steps?
Yes
Determine next steps relevant for your environment and apply
No

### Diagnostics Findings Self-help Flow

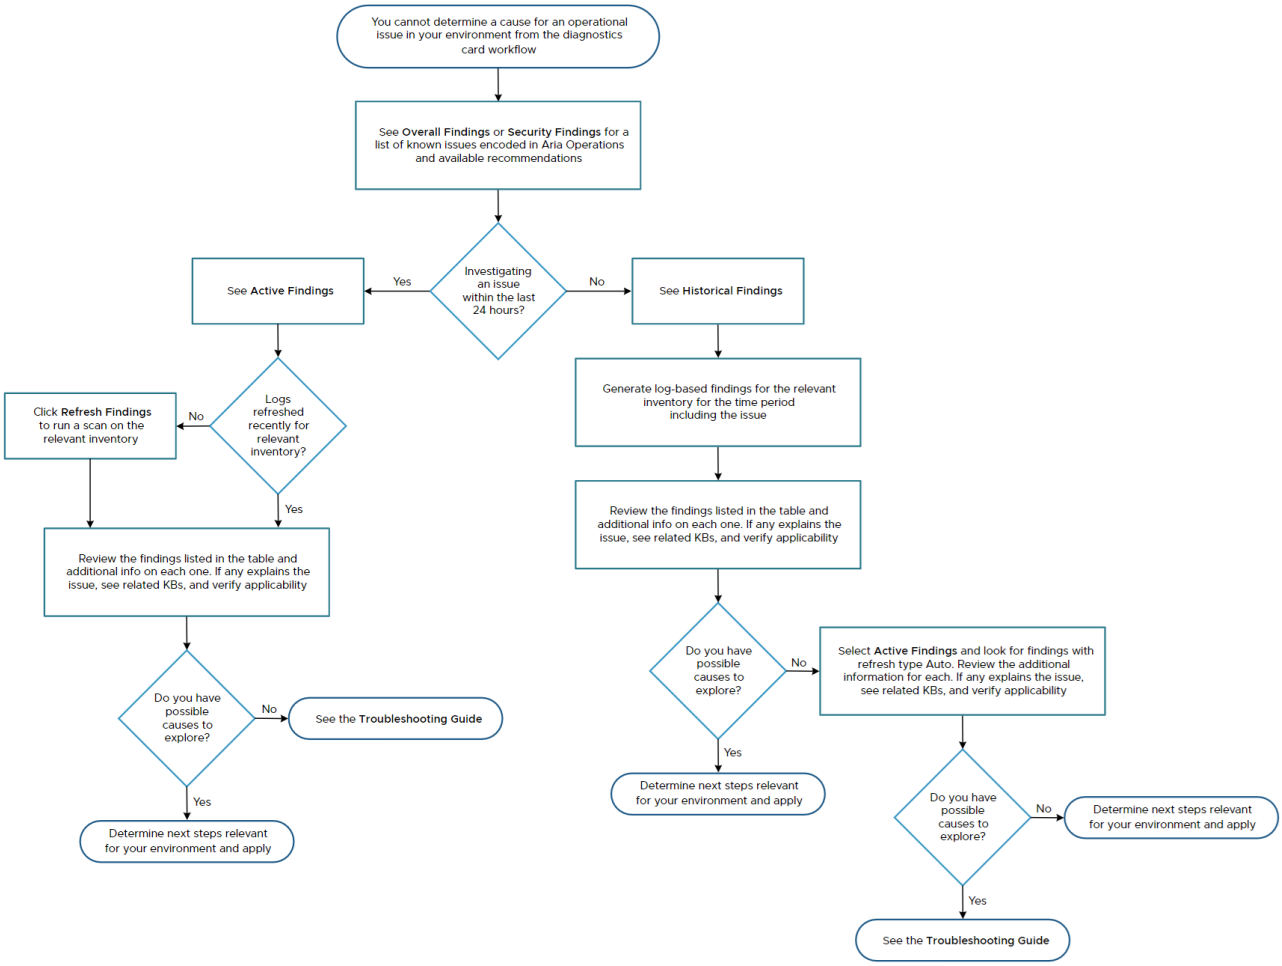


## Diagnostics Rules

Log-based rules detect an actual occurrence of an known issue while property-based rules state that the issue is present in the build and could occur. All these rules relate to specific build numbers. Depending on the findings, you can decide if corrective actions are applicable to your environment and how urgent it is to apply such actions. In some cases, the corrective action is to apply a patch or upgrade, which requires planning, and you cannot immediately apply the recommendations. Diagnostics categorizes findings by severity, component, type, and capability, and reports the number of affected objects for each finding which helps in assessing the impact of findings prioritizing their remediation.

While investigating an issue, logging into a vCenter instance might help you get more details, but is not a required step. In cases when logging into a vCenter instance does not provide sufficient details, refreshing the log-based findings in Diagnostics might provide more relevant data. In case you cannot identify a possible cause for the issue in your environment in the Diagnostics findings page, you can refer to the Troubleshooting Guide, which lists knowledge base articles specific for the conditions that each of the Diagnostics cards reflects and provides steps for resolving such conditions. The Troubleshooting Guide link is available in the Workload Provisioning, vMotion, and Snapshots cards, but includes troubleshooting information for all components that Diagnostics monitors. If the issue relates to a vMotion failure, you can use the log analysis function in VMware Cloud Foundation Operations to retrieve the log statements. You can look up the operation ID on the vMotion detail page.

## What To Do When a Log Scan Fails

When a log scan fails due to long duration for the VMware Cloud Foundation Operations for logs instance, refer to this troubleshooting Knowledge Base article for resolution.

## Setting up Diagnostics for VMware Cloud Foundation

You can activate the diagnostics capabilities for your VMware Cloud Foundation environments by configuring integrations to them, and by installing and configuring VMware Cloud Foundation Operations for logs.

The Diagnostics capability has no license dependency and works on both VMware vSphere Foundation (VVF) and VCF. VMware vSphere Foundation (VVF) provides a subset of the VMware Cloud Foundation (VCF) capabilities and anything that applies to VCF applies to VMware vSphere Foundation (VVF) as well, within the scope of your environment. You can apply only diagnostics capabilities relevant to your setup. For example, if you do not use NSX in your environment, you do not need to take care of the NSX integration. You also do not need to have an SDDC Manager in your environment for Diagnostics to work. The Diagnostic module is installed with the upgrade to VMware Cloud Foundation Operations 8.18.

To fully leverage diagnostics capabilities:

1. Verify that you have deployed and installed VMware Cloud Foundation Operations. See, the Deployment of VMware Aria Operations topic in *Getting Started with VMware Aria Operations Guide*.
2. Install VMware Cloud Foundation Operations for logs. See the Installing VMware Aria Operations for Logs topic in the *Getting Started with VMware Aria Operations Guide*.
3. Integrate VMware Cloud Foundation Operations with VMware Cloud Foundation Operations for logs. See the Configuring VMware Aria Operations for Logs with VMware Aria Operations#undefinedConfiguring VMware Aria Operations for Logs with VMware Aria Operations topic in the *Configuring VMware Aria Operations* guide.
4. Integrate VMware Cloud Foundation Operations for logs with VMware Cloud Foundation Operations to enable communication between VMware Cloud Foundation Operations and the VMware Cloud Foundation Operations for logs instance. See the Using VMware Aria Operations with VMware Aria Operations for Logs topic in the *Administering VMware Aria Operations for Logs* guide.

   **NOTE**
   The diagnostic analysis uses log statements and vSphere events collected by VMware Cloud Foundation Operations for logs. VMware Cloud Foundation Operations can only integrate with a single VMware Cloud Foundation Operations for logs instance. If the required data for a specific product, say a vCenter, is not present in VMware Cloud Foundation Operations for logs, only a partial analysis of that vCenter can be performed. To meet this requirement, you should ensure that all products in the VMware Cloud Foundation Operations inventory report this information to a single instance of VMware Cloud Foundation Operations for logs or deploy multiple instances of VMware Cloud Foundation Operations for logs and configure them to forward the logs and vCenter events to the integrated VMware Cloud Foundation Operations.

5. Activate and configure the following adapters in VMware Cloud Foundation Operations:
   - vCenter. A separate vCenter integration is required for each vCenter that is not part of a VMware Cloud Foundation deployment. See theConfiguring VMware Cloud Foundation Cloud Account in VMware Aria Operations topic in the *Configuring VMware Aria Operations* guide.
     - Enable vSAN collection. See the Verify that the Adapter Instance is Connected and Collecting Data topic in the *Configuring VMware Aria Operations* guide.
   - NSX. A separate NSX integration is required for each NSX instance that is not part of a VMware Cloud Foundation deployment. See the Configuring the NSX Adapter topic in the *Configuring VMware Aria Operations* guide.
   - VMware Cloud Foundation. See the Configuring VMware Cloud Foundation Cloud Account in VMware Aria Operations topic in the *Configuring VMware Aria Operations* guide.
     - Enable NSX and vSAN collections for each discovered vSAN cluster and NSX instance.
   - VMware Aria Automation. See the Configuring VMware Aria Operations for Logs with VMware Aria Operations topic in the *Configuring VMware Aria Operations* guide.
6. Configure vCenter events and ESXi log collection from VMware Cloud Foundation Operations for logs UI. See the Connect VMware Aria Operations for Logs to vSphere Environment topic in the *Administering VMware Aria Operations for Logs* guide.

7. Configure syslog on the configured vCenter to forward its logs to a VMware Cloud Foundation Operations for logs instance. Alternatively, you can deploy the VMware Cloud Foundation Operations for logs agent in vCenter to configure it to send the logs. See the Configure vCenter Server to Forward Log Events to VMware Aria Operations for Logs topic in the *Administering VMware Aria Operations for Logs* guide.
8. Configure log forwarding for SDDC Manager. See Operational Guidance for Intelligent Logging and Analytics.

## Working with VMware Cloud Foundation Diagnostics

You can view a comprehensive summary of your environment on the Diagnostics dashboard.

By navigating to a card, you can view the security and environment findings, state of certificates, operational issues wih vCenter instances, ESXi hosts, and vSAN health, and failures in workload provisioning, vSphere vMotion, and snapshot tasks.

### Monitoring VMware Cloud Foundation Certificates

VMware Aria OperationsVMware Cloud Foundation Operations provides centralised certificate management to monitor and manage all certificates across all VMware Cloud Foundation(VCF) components.

Before you begin monitoring your certificates, you must confirm that the VMware Cloud Foundation integrations for each VCF instance is configured and running. For certificate visibility, the data collection of each integration must be in a heathy state.

To start monitoring your certificates, from the left menu, click **Diagnostics**. On the **Certificates** card, click **View Details** to open the certificates page. The certificates page discovers and displays all the certificates of all VMware Cloud Foundation components.

Use the certificates page to view and export the certificate data. You can click any of the certificates to open the certificate details pane to view in-depth data of the certificate.

Use certificate management to monitor the status and expiration of the following certificate categories:
- Platform Certificate
- Root Certificate
- STS Certificate
- TLS Certificate

If your certificates are about to expire or have already expired, use the **Update** option from the vertical ellipsis to update your certificates.

> **NOTE**
> After you add a new certificate or update your certificate, it takes approximately 10 minutes for the certificate information to be reflected on the certificates page.

### Monitoring vCenter

Diagnostics for VMware Cloud Foundation helps you monitor the operational status and reachability issues of vCenter instances based on analyses of logs and events.

Diagnostics tracks network reachability and service health enabling administrators to proactively identify and address potential problem before they impact operations. The card is linked to a dashboard that reports the reachability of each vCenter and any issues with its services. You can either look at Overall Findings or log in to vCenter instances with reported issues to identify an underlying problem.

For more information, see the vCenter Appliance Availability Dashboard topic in *Configuring VMware Aria Operations* guide.

### Monitoring ESXi Hosts

Use Diagnostics for VMware Cloud Foundation to monitor ESXi hosts that are not in maintenance mode and are reported as not responding state in a vCenter instance.

This report helps you identify hosts that unexpectedly become unresponsive.

The ESXi Host Availability dashboard provides lists and counts of unresponsive, in maintenance mode, powered-off and unknown states of the hosts.

For more information, see the ESXi Host Availability Dashboard topic in the *Configuring VMware Aria Operations* guide.

### Monitoring Workload Provisioning

You can monitor all virtual machine (VM) provisioning operations in your environment using the Workload Provisioning card in the Diagnostics for VMware Cloud Foundation dashboard.

As an administrator, you want to determine:
- What VM operations are running in the environment?
- What is the status of these VM operations. Did they succeed or fail?
- What are the reasons for failure in cases where the workload operations fail?

The Workload Provisioning page lists the total request and failure trend over seven days which enables you to view all VM provisioning operations in your environment, view the VM provisioning status as success or failure, and view the failure reason for each VM provisioing operation. This page reports only the following operations: importing a vApp, creating a VM, cloning to create a new VM, and changing a template into a VM.

#### Troubleshooting Workload Provisioning Issues

To investigate VM provisional operation failures, you can use the trend chart at the top of the page and the filters to gain an understanding of which users are triggering the failures and in which VM, clusters, and vCenter instances. Then, follow the troubleshooting flow provided in What is Diagnostics for VMware Cloud Foundation.

You can also refer to this Knowledge Base article to troubleshoot some common workload provisioning issues.

### Viewing and Resolving Failed Migrations

Diagnostics provides fleet-wide monitoring of successful and failed vSphere vMotion and information that will enable you to diagnose vMotion performance issues and failures.

The vMotion page lists the success and failure trend over seven days, the clusters with the highest number of vMotion over seven days, and every vMotion within the fleet over seven days. The timing parameters for a successful vMotion provide you an understanding of how long the vMotion takes. You can troubleshoot your failures by examining the vSphere vMotion-specific statistics displayed on the dashboard. In addition, the details page provides summary statistics that help you detect unexpected trends.

1. On the vMotion dashboard, click the **virtual machine** name displayed in the table.
2. Click the  icon to open the vSphere vMotion operation and check these three parameters:
   - **Pre-copy Time:** The duration it takes to copy the virtual machine's memory from the source ESXi host to the destination ESXi host. Memory Precopy time is a function of the VM's memory dirty rate on the source ESXi host, the vMotion network transmit rate, and the memory allocation rate on the destination ESX host. vMotion memory precopy is programmed to saturate the vMotion network. In cases where memory pre-copy can't saturate the vMotion network, it likely points to an infrastructure health issue on either the source ESXi, vMotion network, or destination ESXi. Examples of infrastructure health issues are CPU and memory

overcommitment of a source or destination ESXi host and vMotion tasks failing due to network congestion or a malfunctioning NIC on the source or destination host, memory errors, or slow storage.
  - **Pre-copy Bandwidth:** Data transfer rate during the vMotion pre-copy phase.
  - **Switchover Time:** Duration of final memory copy and switchover from source and destination hosts.
3. You can refer to this Knowledge Base (KB) article to troubleshoot some common vSphere vMotion-related issues.

To investigate vMotion failures, you can use the trend chart at the top of the page and the filters to gain an understanding of which users are triggering the failures and in which hosts, clusters, and vCenter instances. You can view the details for the time intervals displayed in the trend chart by using the time filter. Then, follow the troubleshooting flow provided in What is Diagnostics for VMware Cloud Foundation.

**vSphere vMotion Limitations**

- Cross vCenter migrations are not supported.
- Hybrid Cloud Extension (HCX) migrations are not supported.

### Viewing and Resolving Failed Snapshots

Snapshots preserve the state and data of a virtual machine at the time you take the snapshot. You can view the snapshots status using the Snapshots card in the Diagnostics for VMware Cloud Foundation.

Using Diagnostics for VMware Cloud Foundation you can:
- See all snapshot creation events along with their status.
- For each failure view the failure reason.
- View every virtual machine (VM) that requires the consolidation of snapshots. In cases where the file containing snapshot information is incorrect or corrupt, vSphere indicates that a VM requires consolidation. This file does not contain any virtual disk data.

The Findings section displays findings by analyzing the logs for snapshot-related events collected from vCenter instances, ESXi hosts, and VMs.

To investigate snapshot failures, you can use the trend chart at the top of the Snapshots page to gain an understanding of which users are triggering the failures and in which VMs and vCenter instances. Then, follow the troubleshooting flow provided in What is Diagnostics for VMware Cloud Foundation.

**Troubleshooting Snapshot Issues**

You can refer to this Knowledge Base (KB) article to troubleshoot some common snapshot-related issues and best practices.

### Monitoring vSAN Clusters

Diagnostics for VMware Cloud Foundation helps to monitor the vSAN cluster health issues.

The vSAN Clusters card displays the vSAN cluster with issues. The cluster health score ranges from 0 to 100, with 100 being a perfectly healthy cluster based on the conditions of the health findings. A vSAN cluster is considered to have an issue if either the VMware Cloud Foundation Operations generated self-health score for the cluster is below 60, or the cluster has one or more critical alerts.

To view vSAN cluster issues:

1. From the left menu, click **Diagnostics**.
2. On the Diagnostics dashboard, the **vSAN Clusters** card, click **View Dashboard**.

   **NOTE**
   If there are no vSAN cluster issues, **View Dashboard** is deactivated.

3. To identify an issue, click the **Cluster** with red alert.

4. The summary page shows the critical alerts. Click the **critical alert**.
5. Under **Alert Details**, click the **KB article** to troubleshoot the issue.

For more information, see vSAN Health Dashboard.

## Monitoring Objects in Your Managed Environment by Using VMware Aria OperationsVMware Cloud Foundation Operations

Monitoring Objects in Your Managed Environment
Monitoring Objects in Your Managed Environment
You can use VMware Aria OperationsVMware Cloud Foundation Operations to resolve problems that your customers raise, respond to alerts that identify problems before your customers report problems, and generally monitor your environment.

When your customers experience performance problems and call you to resolve the problem, the data that VMware Aria OperationsVMware Cloud Foundation Operations collects and processes is presented to you in graphical forms. You can then compare and contrast objects, understand the relationship between objects, and determine the root cause of problems.

A generated alert notifies you when objects in your environment are experiencing problems. If you resolve the problem based on the alert before your customers notice, then you avoid service interruptions.

You can investigate the problems that generate alerts or that result in calls by using the **Alerts** and **Details** tabs. If you find the root cause of the problem, you might be able to resolve the problem by running an action. The actions change objects in the target system, for example, the VMware vCenter® system, from VMware Aria OperationsVMware Cloud Foundation Operations.

## Enhanced Search Capability

The search function on the top supports locating named objects, dashboards, alerts, and so on, in the system. The search function attempts to match or partially match any string you enter; additional capabilities help you to go swiftly to the item you want. The system presents the item in the Edit context.

### Where you Find Search

The search function appears on all the pages of the VMware Aria OperationsVMware Cloud Foundation Operations in the top menu.

### How Search Works

You start your search by typing in the search bar. VMware Aria OperationsVMware Cloud Foundation Operations displays matching objects types and objects.

The search function supports several common categories you can employ to find the item you seek quickly, as follows:

- Object
- Metric
- Dashboard
- Report
- View
- Alert definition
- Symptom definition
- Recommendation
- Notification
- I.P. Address
- Supermetric

What this means is that in addition to entering a traditional search phrase, for example, a simple string - "VM" - you can also enter one of the listed categories followed by a string or a name. You can then search for objects within the category. For the Object, View and Dashboard categories, the system displays the object in view mode.

If you want quickly to locate a specific dashboard, for example, start typing "dash..." into the search field. The system offers the search term Dashboards. Select the term using the cursor and then enter the dashboard name or part of the name and press Enter. The system finds the dashboard you want, with editing functions available.

Similarly, you can type "alert" or simply "a" in the search field and the system offers Alert Definition. Select the term and enter part of an alert message, for example, "unbalanced." The system returns the alert, "Cluster has an unbalanced workload," presented in the Alert Definition Workspace where you can edit it.

> **NOTE**
> You can type virtual machine in the search bar to list all the virtual machines associated with the host.

**Extend User Search for Alert Assignment**

In VMware Aria OperationsVMware Cloud Foundation Operations now you can search for a user and assign alerts to that user using the following filter options.

- User name
- First Name
- Last Name
- Email Address

VMware Aria OperationsVMware Cloud Foundation Operations displays the details of the searched user in search results, you can use the details to assign the alert to specific user.

**Metric Search**

You can create a search query using a metric, property, object type name, super metric, or instance metric name. Enter the metric name or select the metric, property, or object type from the list of suggestions. The suggestions also display a list of recent searches that match the text you enter in the search bar. You can also delete recent searches. Create a simple query using the metric name and the object type or add conditions and operators to your query to build a more complex query depending on your requirement. Conditions and operators allow you to narrow down the search results. Once you have created the query, click enter to view the search results. The search results are based on the values you select. For more information on how to create queries, see Searching for Metrics, Properties, or Object Types Using Queries.

## Searching for Metrics, Properties, or Object Types Using Queries

The search queries will run across your entire deployment and find all types of objects based on the specified search terms used in a search query. Additionally, VMware Aria OperationsVMware Cloud Foundation Operations provides suggestions to build the search queries, which include recent searches that match the typed text in the global search bar. You can also find specific types of objects using conditions in the search query.

You can create a simple search query starting with a metric, property, or object type, or you can build a more complex query using different conditions to find all types of entities in VMware Aria OperationsVMware Cloud Foundation Operations. Metric queries can be divided into two parts, the metric search query and the metric search result. After you create a metric search query, click **Enter** to view the Metric Search Results page.
You can also use instance metrics and super metrics to search for objects.

1. In the VMware Aria OperationsVMware Cloud Foundation Operations Home page, click the **Search** bar and then select **Metric**.

**NOTE**
When you click the cursor in the search bar, a list of suggestions with the metric and property names and the object types appear. You can select a metric, property, or object type from this list. If you hover over the metric and property names you can view the name of the object types that the metric or property belongs to. If you enter an entity in the search bar, the suggestions change and display options based on your entry. You can start with a metric name and then select the object type it belongs to, or start with an object type and add conditions.

2. Create a simple query with a metric or property name. For example, enter `CPU|Usage %` in the search bar.

   If you start a metric search query with a metric name, you must always mention the object type it belongs to.

   **NOTE**
   If the selected metric belongs to a single object type, then the object type will be selected by default.

3. The object types change based on the metric you enter. Select `Virtual Machine` as the object type.
   A simple metric search query is created as follows: `CPU|Usage % of Virtual Machine`. A green banner with the `Press enter to see the results or continue expanding the query.` message appears.

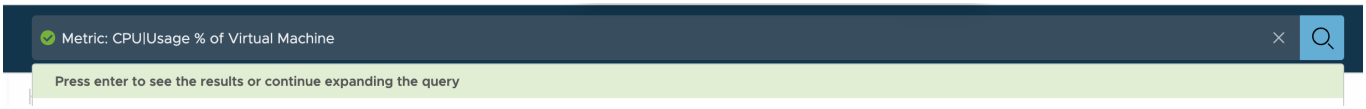


4. Click **Enter** to view the results of this simple search query or continue to expand your search query to narrow down the search results.

   **NOTE**
   If your query is incomplete, has a syntax error, exceeds the character limit of 300, or has more than five different metric names, the search query will not work. In such cases a red banner appears explaining the issue. For example, in case of a syntax error, the `The query contains syntax error(s). Please modify it to get suggestions.` message appears.

5. The metric search query allows you to search for objects based on the conditions you use. Expand the search query with the **where** and **childOf** conditions.
   a) Add a **where** condition to filter the search results. You can combine the where conditions with logical operators to form complex metric search queries.

      **NOTE**
      For string metrics the condition values are case insensitive.

| Operator | Examples |
|---|---|
| **Numeric Operators** | |
| > | `CPU\|Usage % of Virtual Machine where CPU\|Usage % > 15%` |
| >= | `CPU\|Usage % of Virtual Machine where CPU\|Usage % >= 15` |
| < | `CPU\|Usage % of Virtual Machine where CPU\|Usage % < 15` |
| <= | `CPU\|Usage % of Virtual Machine where CPU\|Usage % <= 15%` |

*Table continued on next page*

*Continued from previous page*

| Operator | Examples |
|---|---|
| = | `CPU\|Usage % of Virtual Machine where CPU\|Usage % = 15` |
| != | `CPU\|Usage % of Virtual Machine where CPU\|Usage % != 15` |
| + | `CPU\|Usage % of Virtual Machine where (CPU\|Usage % + Memory\|Usage %) > 15` |
| - | `CPU\|Usage % of Virtual Machine where (CPU\|Usage % - Memory\|Usage %) > 15` |
| * | `CPU\|Usage % of Virtual Machine where (CPU\|Usage % * Memory\|Usage %) > 15` |
| / | `CPU\|Usage % of Virtual Machine where (CPU\|Usage % + Memory\|Usage %)/2 > 15` |
| and | `CPU\|Usage % of Virtual Machine where CPU\|Usage % > 15 and Memory\|Usage % > 15` |
| or | `CPU\|Usage % of Virtual Machine where CPU\|Usage % > 15 or Memory\|Usage % > 15` |
| **String Operators** | |
| equals | `CPU\|Demand MHz of Virtual Machine where Configuration\|Name equals 'Centos'` |
| not equals | `CPU\|Demand MHz of Virtual Machine where Configuration\|Name notEquals 'Centos'` |
| contains | `CPU\|Demand MHz of Virtual Machine where Configuration\|Name contains 'Centos'` |
| notContains | `CPU\|Demand MHz of Virtual Machine where Configuration\|Name notContains 'Centos'` |
| startsWith | `CPU\|Demand MHz of Virtual Machine where Configuration\|Name startsWith 'Centos'` |
| notStartsWith | `CPU\|Demand MHz of Virtual Machine where Configuration\|Name notStartsWith 'Centos'` |

b) Use the **childOf** condition at the end of your query to select an ancestor of the object type to further narrow down your search results. Select the object name, for example, select `vSphere World`.

A complex search query is created as follows: `CPU|Usage % of Virtual Machine where CPU|Usage % > 15 or Memory|Usage % > 15 childOf vSphere World`

6. Click **Enter** to view the metric search results.

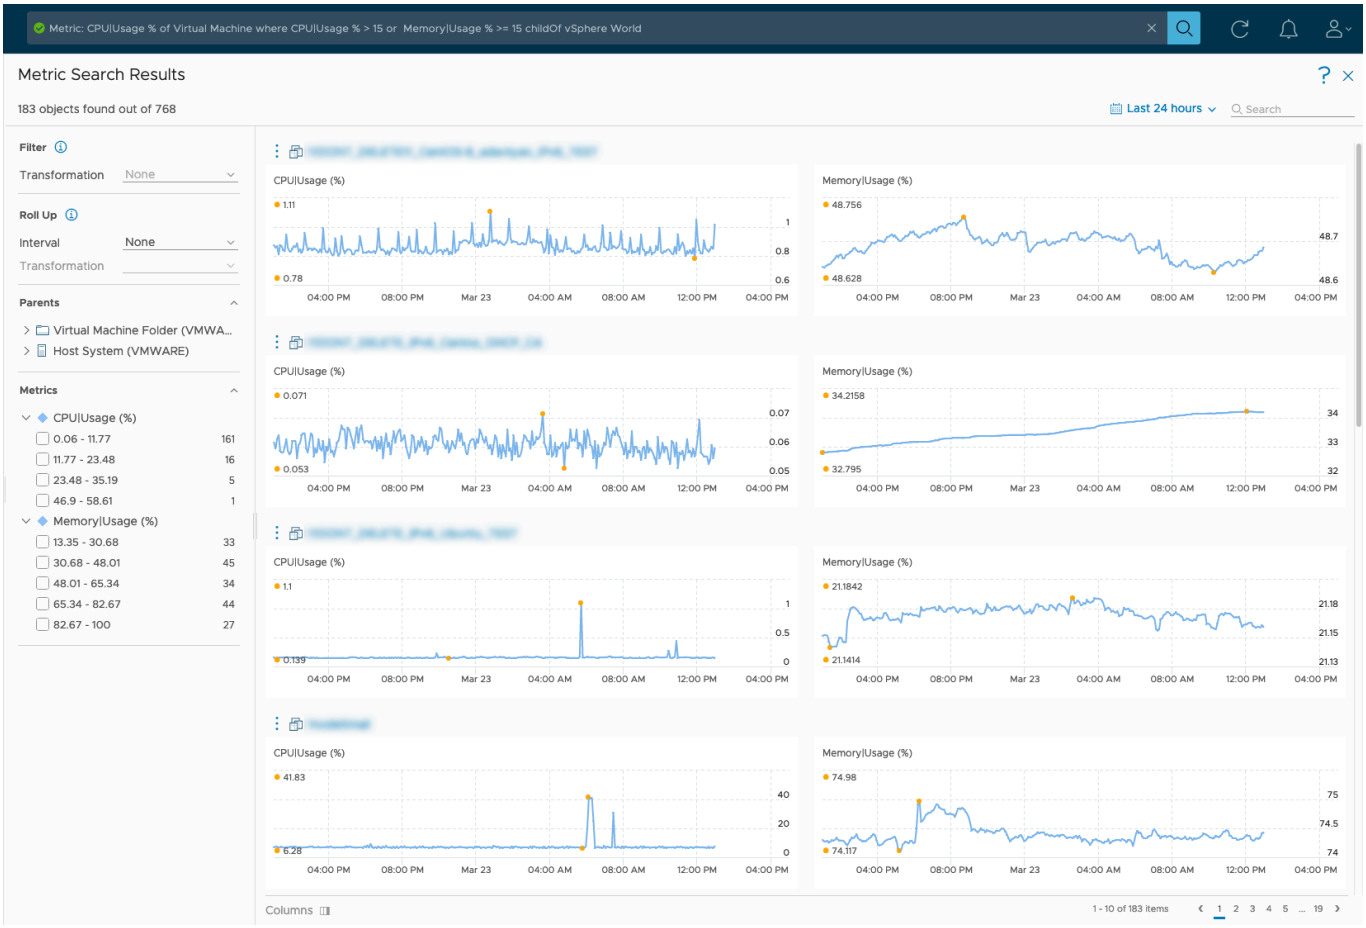


The search result is displayed on the Metric Search Results page.

7. The metric search results page displays the total number of objects that match the query. You can use the filter options on this page to further drill down the results and find the required information.

| Option | Description |
|---|---|
| Filter: Transformation | The Transformation filter applies to the actual query and not the results as it applies to objects before the search query runs and cannot be used on the results.<br><br>Determines what calculation method is applied to the raw data. You can convert the metric values of the search results using the transformation option and also select the type of transformation:<br>• Minimum. The minimum value of the metric over the selected time range.<br>• Maximum. The maximum value of the metric over the selected time range.<br>• Average. The mean of all the metric values over the selected time range. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Sum. The sum of the metric values over the selected time range.<br>• First. The first metric value for the selected time range.<br>• Last. The last value of a metric within the selected time range.<br>• Current. The last available value of a metric if it was last updated not before five collection cycles were complete, otherwise, it is null.<br><br>**NOTE**<br>This option is available only if there is a where condition and the number of objects for the mentioned object type do not exceed the limit of 200 for not more than one month. |
| Roll Up | The roll up filter is used to change data visualization based on the selected time interval and transformation. You can select one of the available options ranging from hour, day, week, month, quarter, or year. For example, if you select Hour as the interval and the Maximum as the transformation, then the system displays maximum values for each one-hour interval. |
| Parents | The parent objects grouped by type for the found objects that the metrics in the query belong to. You can use this option to further narrow down the results. |
| Metric | Displays the metrics value distribution. The metric values are divided into 5 ranges and the number of objects which have the metric lying in that range is available. |
| Time Settings | Use the time settings to select the time range for the query. The number of objects based on metrics in the selected time range is also displayed. You can select a time range to view the information. You can set a time range for a past period or set a future date for the end of the time period. |
| Page Navigation | The page navigation option appears when the results resource count exceeds 10. Use the page navigation to continue viewing search results displayed on the next pages. |
| Columns | If the query contains more than two metrics, you can view the search results of two metrics at a time. Use the columns icon to select which two metrics search results are displayed. |

To further troubleshoot objects, click the vertical ellipsis and select troubleshoot to go to the troubleshooting workbench page. For more information, see Discovering Potential Evidences Using the Troubleshooting Workbench.

## Monitoring VMware Cloud Foundation (VCF) Appliances Health

The VMware Cloud Foundation (VCF) Appliances Health page allows you to view all VCF deployments along with their associated management and workload domains in one place. You can also view the health of the associated management applications such as vCenter, NSX, VMware Aria suite products, and SDDC manager.

- Activate and configure the VMware Cloud Foundation Cloud Account in VMware Aria OperationsVMware Cloud Foundation Operations. For more information, see the 'Configuring VMware Cloud Foundation Cloud Account in VMware Aria OperationsVMware Cloud Foundation Operations' topic in the *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide*.

1. Log in to VMware Aria OperationsVMware Cloud Foundation Operations.
2. From left menu, click **Operations** and then click **VCF Appliances Health**.
   The VCF appliances health page displays the list of VMware Cloud Foundation accounts that you can monitor. The first VMware Cloud Foundation account is selected and the overview of the operations of the selected VMware Cloud Foundation account is displayed by default.
3. You can monitor the VMware Cloud Foundation operations at the account level or the domain level.
   - **Account Level**. Select a VMware Cloud Foundation account that you want to monitor or use the search bar to locate any particular VMware Cloud Foundation account.
   - **Domain Level**. Expand the VMware Cloud Foundation account and click a domain to view the operations of any individual domain.

**Table 320: Monitoring the VMware Cloud Foundation Account**

| Option | Description |
|---|---|
| Availability | Availability indicates if the management application is reachable. |
| Virtual Machines | Virtual Machines indicate the availability of the virtual machines on which the management applications run.<br>**NOTE**<br>The current availability is based on the 'Summary\|Availability' metric value of the virtual machine. For example, VM availability that is less than 50%, is marked as critical, if the criticality of the active symptom definition defined on that metric is set at 50%. The trend is displayed for the last 24 hours only.<br>. |
| Services | Services indicate the availability of all services associated with the application.<br>• Up: Service is available and running. Represented by a green tick icon.<br>• Down: Service is not available. Represented by a red exclamation mark icon.<br>• Unknown: Service status is unavailable. Represented by a grey question mark icon. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | You can monitor the availability of services that constitute an application at the domain level. Expand the application to view the service names, availability, and trend for each of the services. If the services are based on nodes, you can also view the node names that the service runs on.<br><br>You can sort the services based on availability and use the page navigation to view more services.<br><br>**NOTE**<br>The trend is displayed for the last 24 hours only. |
| Passwords | Passwords indicate the validity of passwords associated with the applications. |
| Certificates | Certificates indicate the validity of certificates associated with the applications. |
| Active Alerts | Displays the active alerts of the whole VMware Cloud Foundation account.<br>Click **View Alerts** to view alerts related to the VMware Cloud Foundation object type. The alerts are categorized as critical, immediate, warning, and info.<br><br>At the domain level, you can view the total number of alerts per domain and click **View Alerts** to see alerts related to that domain. |

To view the object hierarchy, you can click the VMware Cloud Foundation account name to view the VMware Cloud Foundation deployment summary tab or click the domain name to view the VMware Cloud Foundation domain objects summary tab. For more information, see the "Summary Tab" topic in the *VMware Aria OperationsVMware Cloud Foundation Operations User Guide*.

## Troubleshooting Workbench Home Page

The **Troubleshooting Workbench** home page is where you find active troubleshooting sessions and recent searches. The active troubleshooting sessions do not persist after you log out from VMware Aria OperationsVMware Cloud Foundation Operations.

**Where You Find the Troubleshooting Workbench Home Page**

- Navigate to the **Troubleshooting Workbench** home page from the left menu by clicking **Operations › Troubleshoot**.

The **Troubleshooting Workbench** home page displays a search bar, a list of active troubleshooting sessions, and recent searches. You can open a session to find potential evidences for your problems.

**How Troubleshooting Workbench Home Page Works**

All troubleshooting workbench sessions that are active in the current login are displayed in the **Active Troubleshooting** section of the **Troubleshooting Workbench** home page. Changes that you make to the scope, time, or potential evidences in the troubleshooting workbench page are not be saved on logging out. The next time you log in to VMware

Aria OperationsVMware Cloud Foundation Operations, the sessions that were earlier under **Active Troubleshooting** are displayed under **Recent Searches**.

### Discovering Potential Evidences Using the Troubleshooting Workbench

The Troubleshooting Workbench is where you perform advanced troubleshooting tasks on an alert that triggered on an object. You can investigate both known and unknown issues in VMware Aria OperationsVMware Cloud Foundation Operations.

#### Where You Find the Troubleshooting Workbench

You can start the Troubleshooting Workbench with an alert in context from the alert information page, or you can search for an object or metric and start the Troubleshooting Workbench to investigate known or unknown issues related to the object.

- To start the Troubleshooting Workbench with an alert in context, in the menu, click **Operations › Troubleshoot › Alerts**. Click an alert from the alert list and click the **TROUBLESHOOT** button on the top right.
- To start the Troubleshooting Workbench with an alert in context, in the menu, select a group, custom data center, application, or inventory object from the **Inventory** in the left menu. Click the object and then the **Alerts** tab. Cick the **TROUBLESHOOT** button on the top right.
- On the Troubleshooting Workbench page (**Operations › Troubleshoot**), to investigate known or unknown issues with an object in context, search for the object.
- On the Troubleshooting Workbench page (**Operations › Troubleshoot**), search for a metric and start a troubleshooting session in the context of the metric. The metric search has the same enhanced search capabilities that is available in the global search. For more information about the advanced search capabilities, see the topic, Enhanced Search Capability. Recent searches are saved and displayed when you start a search. After you locate the metric you are looking for, VMware Cloud Foundation Operations displays the metrics within the Troubleshooting Workbench page. You can minimize the search result and come back to it later. If you are done seeing the metrics, you can close the search result window. To start troubleshooting using the metric, click the vertical ellipsis and then **Troubleshoot**.

#### How the Troubleshooting Workbench Works

You look for potential evidences of a problem within a specific scope and time range. The **Selected Scope** control on the left of the Troubleshooting Workbench page is where you vary the scope. You can vary the scope in the following ways:

- You can select only the object that you are investigating, or include several upstream and downstream relationships by increasing the scope. As you increase the scope, more objects are displayed in the inventory tree.
- You can select a custom scope to include objects of your choice. Click **Custom** to open an interactive window where you use the pointer to visually rearrange your objects, view relationships and add peers to modify the relationships. To see details about the object, place the pointer for a few seconds above the object. You can reset a custom scope to start all over again.
- You can use the drop-down menu to narrow down the type of objects displayed.

The default time range is two hours, and thirty minutes before the alert triggered when the context is alert based, or one hour before the current time, when the context is object based. You can select a different time range, up to seven days, using the date and time controls.

The potential evidences are based on Events, Property Changes, and Anomalous Metrics which are displayed on the right of the Troubleshooting Workbench change in the **Potential Evidence** tab. Information in these sections is displayed as cards.

**Events**

Displays events, based on a change in the metrics. Events for metrics that have breached the usual behavior, and major events that have occurred within the selected scope and time are displayed. The cards are based on dynamic thresholds for a metric, which is calculated based on historical and incoming data.

**Property Changes**

Displays important configuration changes that occurred within the selected scope and time. Both single and multiple property changes are displayed. For multiple property changes, you can view the latest and previous changes.

**Anomalous Metrics**

Metrics which have shown drastic changes within the selected scope and time. Ranks the results based on the degree of change. The most recent anomalous metric based on a time-sliced comparison in the current time range is given the highest weightage.

You can explore more details about any of the cards displayed in the Troubleshooting Workbench by clicking the card pop-out option. You can close a card and it is no longer displayed in the Troubleshooting Workbench. To load the cards again, click **Go** in the **Time Range**.

When you pin a metric, it appears in the **Metrics** tab of the Troubleshooting Workbench. You can perform further investigation on the metric in the Metrics tab. You can compared the pinned metrics with other metrics displayed in the tab. You can close the pinned metrics and browse other metrics for specific objects.

Similarly, the **Alerts** and **Events** tabs are where you investigate the potential evidences further. You can filter and group alerts. If you want to focus on the alerts for a specific object in your selected scope, you can clear all the alerts and then click the object in the scope. You can view logs in the **Logs** tab, and launch VMware Aria Operations for Logs for further investigation of the logs.

## Monitoring and Responding to Alerts

Alerts indicate a problem in your environment. Alerts are generated when the collected data for an object is compared to alert definitions for that object type and the defined symptoms are true. When an alert is generated, you are presented with the triggering symptoms, so that you can evaluate the object in your environment, and with recommendations for how to resolve the alert.

Alerts notify you when an object or group of objects are exhibiting symptoms that are unfavorable for your environment. By monitoring and responding to alerts, you stay aware of problems and can react to them in a timely fashion.

Generated alerts drive the status of the top-level badges, Health, Risk, and Efficiency.

In addition to responding to alerts, you can generally respond to the status of badges for objects in your environment.

You can take ownership of an alert or assign alerts to other VMware Aria OperationsVMware Cloud Foundation Operations users.

### Monitoring Alerts in VMware Aria OperationsVMware Cloud Foundation Operations

Monitoring Alerts
Monitoring Alerts
You can monitor your environment for generated alerts in several areas in VMware Aria OperationsVMware Cloud Foundation Operations. The alerts are generated when the symptoms in the alert definition are triggered, letting you know when the objects in your environment are not operating within the parameters you defined as acceptable.

Generated alerts appear in many areas of VMware Aria OperationsVMware Cloud Foundation Operations so that you can monitor and respond to problems in your environment.

**Alerts**

Alerts are classified as Health, Risk, or Efficiency. Health alerts indicate problems that require immediate attention. Risk alerts indicate problems that must be addressed shortly, before the problems become immediate health problems.

Efficiency alerts indicate areas where you can reclaim wasted space or improve the performance of objects in your environment.

You can monitor the alerts for your environment in the following locations.

- Alerts
- Health
- Risk
- Efficiency

You can monitor alerts for a selected object in the following locations.

- Alert Details, including the **Summary**, **Timeline**, and **Metric Charts** tabs
- **Summary** tab
- **Alerts** tab
- **Events** tab
- Custom dashboards
- Alert notifications

## Working with Alerts

Alerts indicate a problem that must be resolved so that triggering conditions no longer exist and the alert is canceled. Suggested resolutions are provided as recommendations so that you can approach the problem with solutions.

As you monitor alerts, you can take ownership, suspend, or manually cancel alerts.

When you cancel an alert, the alert and any symptoms of type message event, or metric event are canceled. You cannot manually cancel other types of symptoms. If a message event symptom or metric event symptom triggered the event, then the alert is effectively canceled. If a metric symptom or property symptom triggered the alert, a new alert might be created for the same conditions in the next few minutes.

The correct way to remove an alert is to address the underlying conditions that triggered the symptoms and generated the alert.

## Migrated Alerts

If you migrated alerts from a previous version of VMware Aria OperationsVMware Cloud Foundation Operations, the alerts are listed in the overview with a canceled status, but alert details are not available.

### User Scenario: Monitor and Process Alerts in VMware Cloud Foundation OperationsVMware Cloud Foundation Operations

User Scenario: Monitor and Process Alerts
User Scenario: Monitor and Process Alerts
Alerts in VMware Cloud Foundation OperationsVMware Cloud Foundation Operations notify you when objects in your environment have a problem. This scenario illustrates one way that you can monitor and process alerts for the objects you are responsible for.

An alert is generated when one or more of the alert symptoms are triggered. Depending on how the alert is configured, the alert is generated when one symptom is triggered or when all the symptoms are triggered.
As the alerts are generated, you must process the alerts based on the negative effect they have on objects in your environment. To do the processing, you start with Health alerts, and process them based on criticality.
As a virtual infrastructure administrator, you review the alerts at least twice a day. As part of your evaluation process in this scenario, you encounter the following alerts:

- Virtual machine has unexpected high CPU workload.
- Host has a memory contention that a few virtual machines cause.
- Cluster has many virtual machines that have a memory contention because of memory compression, ballooning, or swapping.

1. From the left menu, click **Operations › Alerts**.
2. Select **Time** in the Group By filter and the click the down arrow in the Created On column, so the most recent alerts are listed first .
3. In All Filters, select **Criticality > Warning**
   You have listed all the Warning alerts in order of when they fired, with the most recent alerts appearing first.
4. Review the alerts by name, the object on which it was triggered, the object type, and the time at which the alert was generated.
   For example, do you recognize any of the objects as objects that you are responsible for managing? Do you know that the fix that you will implement in the next hour will fix any of the alerts that are affecting the Health status of the object? Do you know that some of your alerts cannot be resolved currently because of resource constraints?
5. To indicate to other administrators or engineers that you are taking ownership of the `Virtual machine has unexpected high CPU workload` alerts, click the selected alerts, click **Actions** on the menu bar, and click **Take Ownership**.
   The Assigned to: field in Alert Details updates with your user name.
6. To assign the ownership of the `Virtual machine has unexpected high CPU workload` alert to another user, click the alert, click **Actions** on the menu bar, and click **Assign to**.
7. Enter the name of user to whom you want to assign the ownership of the alert and click **Save**.
   The Assigned to: field in Alert Details updates with the name of the user you have assigned the alert to.

   **NOTE**
   You can remove the ownership assigned to a user by clicking the alert and selecting the **Release Ownership** option from the **Actions** menu.

8. To take ownership and temporarily exclude the alert from affecting the state of the object, select the `Host has memory contention caused by a few virtual machines` alert in the list. Then click **Actions** on the menu bar and click **Suspend**.
   a) To suspend the alert for an hour, enter `60`.
   b) Click **OK**.

   The alert is suspended for 60 minutes and you are listed as the owner in the alert list. If it is not resolved in an hour, it returns to an active state.
9. Select the row that contains the `Cluster has many Virtual Machines that have memory contention due to memory compression, ballooning or swapping` alert. Then click **Actions** on the menu bar and click **Cancel Alert** to remove the alert from the list.

   This alert is a known problem that you cannot resolve until the new hardware arrives.

   The alert is removed from the alert list, but this action does not resolve the underlying condition. The symptoms in this alert are based on metrics, so the alert will be generated during the next collection and analysis cycle. This pattern continues until you resolve the underlying hardware and workload distribution issues.

You processed the critical health alerts and took ownership of the ones to resolve or troubleshoot further.

Respond to an alert. See User Scenario: Respond to an Alert in the Health Alert List.

**User Scenario: Respond to an Alert in the Health Alert List**

Generated alerts in VMware Aria OperationsVMware Cloud Foundation Operations appear in the alert lists. You use the alert lists to investigate, resolve, and begin troubleshooting problems in your environment.

- Process and take ownership of the alerts you troubleshoot and resolve. See User Scenario: Monitor and Process Alerts in .

- Review information about how the Power Off Allowed setting works when you run actions. See the section Working with Actions That Use Power Off Allowed in the VMware Aria OperationsVMware Cloud Foundation Operations Information Center.
- Process and take ownership of the alerts you troubleshoot and resolve. See User Scenario: Monitor and Process Alerts in .
- Review information about how the Power Off Allowed setting works when you run actions. See Working with Actions That Use Power Off Allowed section in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* .

In this scenario, you investigate and resolve the `Virtual machine has an unexpected high CPU workload` alert. The alert might be generated for more than one virtual machine.

1. From the left menu, click **Operations › Alerts**.
2. To limit the list to virtual machine alerts, click **All Filters** on the toolbar.
   a) Select **Object Type** in the drop-down menu.
   b) Enter `virtual machine` in the text box.
   c) Click **Enter**.

   The alerts list displays only alerts based on virtual machines.
3. To locate the alerts by name, enter `high CPU workload` in the **Quick filter (Alert)** text box.
4. In the list, click the **Virtual machine has an unexpected high CPU workload** alert name.
5. Review the information. To show the recommendations, click **Configuration > Recommendations** in the left pane .

| Option | Evaluation Process |
|---|---|
| Alert Description | Review the description so that you better understand the alert. |
| Recommendations | Do you think that implementing one or more of the recommendations can resolve the alert? |
| What is Causing the Issue? | Do the triggered symptoms support the recommendations? Do the other triggered symptoms contradict the recommendation, indicating that you must investigate further?<br><br>In this example, the triggered symptoms indicate that the virtual machine CPU demand is at a critical level and that the virtual machine anomaly is starting to get high. |
| Non-Triggered Symptoms | Some alerts are generated only when all the symptoms are triggered. Others are configured to generate an alert when any one of the symptoms are triggered. If you have non-triggered symptoms, evaluate them in the context of the triggered alerts.<br><br>Do the non-triggered symptoms support the recommendations? Do the non-triggered symptoms indicate that recommendations are not valid and that you must investigate further? |

6. To resolve the alert based on the recommendation to check the guest applications to determine whether a high CPU workload is an expected behavior, click the **Action** menu on the center pane toolbar and select **Open Virtual Machine in vSphere Client**.
   a) Log in to the vCenter instance using your vSphere credentials.
   b) Start the console for the virtual machine and identify which guest applications are consuming CPU resources.

7. To resolve the alert based on the recommendation to add more CPU capacity to this virtual machine, click **Set CPU Count for VM**.
   a) Enter a new value in the **New CPU** text box.

      The value that appears is the calculated suggested size. If VMware Aria OperationsVMware Cloud Foundation Operations was monitoring the virtual machine for six or more hours, depending on your environment, the value that appears is the CPU recommended size metric.

   b) To allow power off or to create a snapshot, depending on how your virtual machines are configured, select the following options.

| Option | Description |
|---|---|
| Power Off Allowed | Shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without any regard for the state of the operating system.<br><br>In addition to the question whether the action shuts down or powers off a virtual machine, you must also consider whether the object is powered on and what settings are applied. |
| Snapshot | Creates a snapshot of the virtual machine before you add CPUs.<br><br>If the CPU is changed with CPU Hot Plug activated, then the snapshot is taken with the virtual machine running, which consumes more disk space. |

   c) Click **OK**.

   The action adds the suggested number of CPUs to the target virtual machine.
8. Allow several collection cycles to run after implementing the suggested changes and check the alert list.

If the alert does not reappear after several collection cycles, it is resolved. If it reappears, further troubleshooting is required.

## Monitoring and Responding to Problems

The organization of the tabs and options in VMware Aria OperationsVMware Cloud Foundation Operations provides a built-in workflow that you can use when you work with objects in your environment.

The tabs, **Summary**, **Metrics**, **Alerts**, **Topology**, **Capacity**, and so on, provide progressive levels of detail about the selected object. As you work through the tabs, starting with the high level **Summary** and **Alerts** tabs, you see the general state of an object. The **Details** tab is specific to data views and lists.

As you monitor objects in your environment, you discover which tabs provide the information that you need when you are investigating problems.

### Inventory Page and Inventory Detailed View

You can view details of vCenter Servers from the **Inventory** page. You can also view a detailed list of the inventory in your VMware Aria Operations environment where the objects are grouped logically. When you select an object from the detailed view, you see information such as summary information, metrics, alerts, capacity, compliance, and other details specific to the selected object.

#### Inventory Page

To access the **Inventory** page, select **Inventory** from the left panel.

> **NOTE**
> If you see the Inventory detailed view, click **Basic View** at the top in the **Inventory** panel.

The **Inventory** page lists the vCenter Servers that have been configured. At the top of the page, you can view environment details such as the number of VCF accounts, vCenter accounts, data centers, clusters, hosts, VMs, and datastores.

**Table 321: Options in the Data Grid of the Inventory Page**

| Option | Description |
|---|---|
| vCenter Name | Displays the names of the vCenter Server that are configured. |
| Type | Displays all account types. |
| Data centers | Displays the number of data centers in the environment. |
| vSphere Clusters | Displays the number of vSphere clusters that have been configured. |
| Hosts | Displays the number of hosts in the environment. |
| Virtual Machines | Displays the number of VMs in the environment. |
| Datastores | Displays the number of datastores in the environment. |

#### Inventory Detailed View

The Inventory detailed view is a panel that displays different types of logical grouping/hierarchies for the objects in the VMware Aria Operations environment. You can view either a tree or list view and then select the object to view object related information.

To access the **Inventory Detailed View**, from the left panel, click **Inventory**, and then click on a vCenter Server instance from the data grid or click **Detailed View** from the top. You will see the **Inventory** panel with the Inventory detailed view.

**Table 322: Options from the Inventory Detailed View**

| Option | Description |
|---|---|
| Tree view and List view | Toggle between the options to view the inventory in either a tree view or a list view. When you select a certain object your view is scoped around that object.<br>In the list view, you can see all the descendant and ancestor objects of the selected object. When you double click on any of the objects in the list view, you change the scope and the new object becomes the primary object of the view. You can select the number of ancestors or |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | descendants for the object you want to view from the **Levels** option. |
| vSphere Hosts and Clusters | Select a vSphere host or cluster object to view the object details. |
| vSphere Storage | Select a vSphere Storage object to view the object details. |
| vSphere Networking | Select a vSphere Networking object to view the object details. |
| Custom Group and Datacenters | Select an object from Custom Groups, Custom Datacenters, or Business Applications to view the object details. |
| Integrations | Select an object from Others and All Objects to view the object details. |

**NOTE**
If there are more than ten objects under a section in the Inventory detailed view, you can search for an object using the Search option. If there are more than thousand objects in a section, use the **View More** button under the last object that is displayed, to view the rest of the objects. Use the **Type to see more** button to find objects if there are more than two thousand objects.

### Evaluating Object Information Using Badge Alerts and the Summary Tab

The Summary tab that is associated with the other object tabs summarizes Heath, Risk, and Efficiency badge alerts for the selected object and displays the top alerts that lead to the current state.

Use this tab as an overview of alerts for an object, object group, or application - to evaluate the effect that alerts are having on an object and to begin troubleshooting problems. For more detail on the badge Alerts, click **Badge Alerts**, further to the right on the tool bar.

#### Badge Alert Types

The Health, Risk, and Efficiency badge states are based on the number and criticality of the generated alerts for the selected object.

- Health alerts indicate problems that affect the health of your environment and require immediate attention to ensure that service to your customers is not affected.
- Risk alerts indicate problems that are not immediate threats but must be addressed shortly.
- Efficiency alerts tell you where you can improve performance or reclaim resources.

#### Alerts for an Object or an Object Group

For a single object, the Top alerts are the alerts generated for the object. The Top Alerts for Children are the alerts generated for any child or other descendant objects in the currently selected navigation hierarchy. For example, if you are working with a host object in the vSphere Host and Clusters navigation hierarchy, children can include virtual machines and datastores.

Object groups can include one object type, such as hosts, or multiple objects types, such as hosts, virtual machines, and datastores. When you are working with object groups, all the group member objects are children of the group container. The most critical generated alerts for the member objects appear as Top Alerts for Children.

For an object group, the only Top Alerts that might be generated are the predefined group population alerts. If the average health is above the Warning, Immediate, or Critical threshold, a group population alert considers the health of all group members and is triggered. If a group population alert is generated, the alert affects the badge score and color. If a group population alert is not generated, then the badges are green. This behavior is because an object group is a container for other objects.

### Summary Tab and Related Hierarchies

The alerts that appear on the **Summary** tab for an object can vary depending on the currently selected hierarchy in the Related Hierarchies in the left pane.

Depending on the selected hierarchy, you see different alerts and relationships on the **Summary** tab for an object. The current focus object name is on the center pane title bar, but the children alerts depend on the relationships that the highlighted hierarchy defined in the Related Hierarchies list in the upper left pane. For example, if you are working with a host object relative to virtual machines in the vSphere Hosts and Clusters hierarchy, then children commonly include virtual machines and datastores. But if you are working with the same host as a member of an object group, then any alerts on virtual machines that are also members of the group do not appear. The alerts do not appear because the host and the virtual machines are considered children of the group and peers among each other. In this example, the focus of the **Summary** tab is the host in the context of the group, not the vSphere Hosts and Clusters hierarchy.

### Summary Tab Evaluation Techniques

You can evaluate the state of objects, starting with the **Summary** tab, by using one or more of the following techniques.

- Select an object or object group, click the alerts on the **Summary** tab, and resolve the problems that the alert indicates.
- Select an object, review the alerts on the **Summary > Alerts** tab, and select other objects, comparing the volume and types of alerts generated for different objects.

## Summary Tab

The Summary tab provides an overview of the state of the selected object, group, or application. Use this tab to evaluate the impact that alerts are having on the object and use the information to begin troubleshooting problems.

### How the Summary Tab Works

Based on the object selected, the following summary tabs are displayed:

- VM Summary Tab
- Datastore Summary Tab
- Host Summary Tab
- Cluster Summary Tab
- Custom Group and Container Summary Tab

### Where to Find the Summary Tab

- In the menu, click **Inventory**, then select a group, custom data center, application, or inventory object.

**Understanding the Summary Tab**

**Table 323: Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. It also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Consumer | Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, CPU, and Memory. |
| Provider | Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage. |
| Cluster | Displays the cluster details of the selected object. |
| Datastore | Displays the datastore details of the selected object. |

**Datastore Summary Tab**

The Datastore Summary tab provides an overview of the state of the selected datastore. For the selected object, the Datastore Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the datastore and use the information to begin troubleshooting problems.

Understanding the Datastore Summary Tab

**Table 324: Datastore Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Utilization | This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available. |
| Performance | This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart. |
| Configuration | This widget displays the configuration details for the selected datastore object. |

### Host Summary Tab

The Host Summary tab provides an overview of the state of the selected host. For the selected object, the Host Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the host and use the information to begin troubleshooting problems.

Understanding the Host Summary Tab

**Table 325: Host Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Utilization | This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available. |
| Performance | This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart. |
| Configuration | This widget displays the hardware, CPU, and Network configuration details of the host. |
| Ping Statistics | This widget displays the availability of end points that exist in your environment and other details. For more information, see Configuring Ping Adapter Instances. |

### VM Summary Tab

The VM Summary tab provides an overview of the state of the selected VM. For the selected object, the VM Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the VM and use the information to begin troubleshooting problems.

Understanding the VM Summary Tab

**Table 326: VM Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Utilization | This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available. |
| Performance | This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart. |
| Configuration | This widget displays the virtual hardware, resource allocation, tools, and Network configuration details of the virtual machine. |

### Cluster Summary Tab

The Cluster Summary tab provides an overview of the state of the selected cluster. For the selected object, the Cluster Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the cluster and use the information to begin troubleshooting problems.

Understanding the Cluster Summary Tab

**Table 327: Cluster Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Virtual Machine Remaining | This widget displays the remaining virtual machines in the cluster. To see the details of the remaining virtual machines, click the Virtual Machine Remaining card. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Utilization | This widget is used to find out the trends in capacity used by a selected datastore as against the total capacity available. |
| Performance | This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart. |
| Configuration | This widget displays the configuration details of the cluster. |
| GPU Information | This widget is only visible if the cluster has GPU configured hosts. It displays the following GPU information:<br>• Number of GPUs: The total number of GPUs in the cluster.<br>• Compute Utilization: The compute utilization percentage of a GPU<br>• Memory Usage: The percentage of memory currently in use out of the total available memory. |
| Metadata | This widget displays the metadata details of the cluster. |

### vCenter Server and Data Center Summary Tab

The vCenter Server and data center Summary tab provides an overview of the state of the selected data center or vCenter. For the selected object, the vCenter server or data center Summary tab displays the alerts as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the vCenter server or data center and use the information to begin troubleshooting problems.

Understanding the vCenter Server and Data Center Summary Tab

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br>To see the alerts for the object, click the labels of the alert. |
| Consumer | Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, CPU, and Memory. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Provider | Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage. |
| vSphere Distributed Switch Name | Displays the details of the vSphere distributed switch. |
| Metadata | Displays the metadata details of the data center. |
| Cluster | Displays the cluster details of the selected object. |
| Datastore | Displays the datastore details of the selected object. |

### Resource Pool Summary Tab

The Resource Pool Summary tab provides an overview of the state of the resources in the resource pool. For the selected resource, the Resource Pool Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the resource pool and use the information to begin troubleshooting problems.

Understanding the Resource Pool Summary Tab

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br>To see the alerts for the object, click the labels of the alert. |
| Utilization | This widget is used to find out the trends in capacity used by the selected resource pool as against the total capacity available. |
| Performance | This widget displays the summary metrics about the overall performance of the object. Click each metric to see the expanded chart. |
| Resource Pool | This widget lists the resource pool name, cpu status, and memory status of the resources that are part of the corresponding resource pool. |

### Custom Group and Container Summary Tab

The Custom Group and Container Summary tab provides an overview of the state of the selected group or a container. For the selected object, the Custom Group and Container Summary tab displays the alerts and metrics as they affect the health, risk, or efficiency. Use this tab to evaluate the impact that alerts are having on the group or a container and use the information to troubleshoot the problems.

Understanding the Custom Group and Container Summary Tab

**Table 328: Custom Group and Container Summary Tab Options**

| Option | Description |
|---|---|
| Recommended Actions | This widget displays the health status for the selected object and its descendants. It also displays recommendations to solve problems in an instance.<br><br>The badges provide a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the badge. |

### Workload Management Activated Cluster Summary Tab

The Workload Management activated cluster is a cluster with Kubernetes activated, running on vSphere (also called Supervisor cluster). It hosts a type of resource pool called Namespaces. The Workload Management Enabled Cluster Summary tab provides an overview of the state of the selected cluster.

Understanding the Cluster Summary Tab

**Table 329: Workload Management Enabled Cluster Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object and whether the Workload Management is activated or deactivated. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the badge. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Virtual Machine Remaining | The virtual machine remaining number is based on the average profile. The virtual machine remaining numbers are calculated when you activate one or more custom profiles from the policy. The overall virtual machine remaining is based on the most constrained profile. |
| Utilization | This widget is used to find out the trends in capacity used by a selected cluster as against the total capacity available.<br>The key utilization indicators are:<br>• CPU Capacity Usage<br>• Memory Usage<br>• Memory Balloon<br>• Disk Total IOPS<br>• Disk Total Throughput<br>• Network Usage Rate |
| Performance | This widget displays the summary metrics about the overall performance of the object. It displays the latest value and a trend line of the various key performance indicators in a color that indicates its health based on the symptom associated with the metrics. Click each metric to see the expanded chart.<br><br>The key performance indicators are:<br>• Max VM Memory Contention<br>• Worst Consumer Disk Latency<br>• Consumers with Memory Contention<br>• Consumers with CPU Ready<br>• Physical Network Packets Dropped<br>• Virtual Network Packets Dropped |
| Configuration | This widget displays the hardware, CPU, and Network configuration details of the host. |
| Namespaces | Lists the configuration status, current version and Kubernetes status of the namespaces in the cluster. |

### Namespace Summary Tab

A namespace sets the resource boundaries where vSphere Pods and Tanzu Kubernetes clusters created by using the Tanzu Kubernetes Grid Service can run. The Namespace summary tab provides an overview of the state of the selected Namespace.

Understanding the Namespace Summary Tab

**Table 330: Namespace Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status based on the alert type.<br><br>To see the alerts for the object, click the badge . |
| Utilization | This widget is used to find out the trends in capacity used by a selected namespace as against the total capacity available.<br>The key utilization indicators are:<br>• CPU Usage<br>• Consumed Memory |
| Performance | This widget displays the summary metrics about the overall performance of the object. It displays the latest value and a trend line of the various key performance indicators in a color that indicates its health based on the symptom associated with the metrics. Click each metric to see the expanded chart.<br><br>The key performance indicators are:<br>• Worst Consumer CPU Ready<br>• Worst Consumer Memory Contention<br>• Consumers with Memory Contention<br>• Consumers with CPU Ready |
| Configuration | This widget displays the following configuration details about the Namespaces:<br>• Configuration status<br>• Virtual Machines<br>• Number of Tanzu Kubernetes clusters<br>• Pods |

### vSphere Pod Summary Tab

vSphere Pods run containers without needing to customize a Kubernetes cluster. You can deploy vSphere Pods directly on ESXi hosts It hosts a type of resource pool called Namespace. The vSphere Pod Summary tab provides an overview of the state of the vSphere Pods.

Understanding the vSphere Pod Summary Tab

**Table 331: vSphere Pod Tab Summary Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining until the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Utilization | This widget is used to find out the trends in capacity used by a selected vSphere Pod as against the total capacity available.<br>The key utilization indicators are:<br>• CPU Usage<br>• Free Memory<br>• Guest Page in Rate per second<br>• Virtual Disk Total IOPS<br>• Virtual Disk Total Throughput |
| Performance | This widget displays the summary metrics about the overall performance of the object. It displays the latest value and a trend line of the various key performance indicators in a color that indicates its health based on the symptom associated with the metrics. Click each metric to see the expanded chart.<br><br>The key performance indicators are:<br>• CPU Queue<br>• Disk Queue<br>• CPU Ready |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • CPU Co-stop<br>• Memory Contention<br>• Virtual Disk Total Latency<br>• Network Transmitted Packets Dropped |
| Configuration | This widget displays the hardware, CPU, and Network configuration details of the host. |

### Tanzu Kubernetes cluster Summary Tab

The Tanzu Kubernetes cluster runs Kubernetes workloads natively on the hypervisor layer. The Tanzu Kubernetes cluster Summary tab provides an overview of the state of the Tanzu Kubernetes clusters.

Understanding the Tanzu Kubernetes cluster Summary Tab

**Table 332: Tanzu Kubernetes cluster Tab Summary Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the badge. |
| Utilization | This widget is used to find out the trends in capacity used by a selected Tanzu Kubernetes cluster as against the total capacity available.<br>The key utilization indicators are:<br>• CPU Usage<br>• Consumed Memory |
| Performance | This widget displays the summary metrics about the overall performance of the object. It displays the latest value and a trend line of the various key performance indicators in a color that indicates its health based on the symptom associated with the metrics. Click each metric to see the expanded chart.<br><br>Key performance indicators are:<br>• Worst Consumer CPU Ready |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Worst Consumer Memory Contention<br>• Consumers with Memory Contention<br>• Consumers with CPU Ready |

**Summary Tabs for VMware Cloud Foundation**

The summary tab of VMware Cloud Foundation (VCF) provides account details, overall SDDC inventory of the the organization and their the key components, aggregated metrics, alerts, and so on.

## VCF Summary Tab

The VCF Summary tab provides organization details, the overall inventory of the organization including key components, alerts, and so on.

Where To View VCF Summary

From the left menu, click **Data Sources › Integrations › Accounts**. Click VMware Cloud Foundation, click the vertical ellipsis against the VCF account, and then select **Object Details**.

**Table 333: VCF Summary Options**

| Option | Description |
|---|---|
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts indicate issues that require immediate attention.<br>• Warning alerts indicate that you must look into issues shortly.<br>• Info alerts indicate that you can reclaim resources. |
| Consumer | Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, vCPU, RAM, and Provision. |
| Provider (Usable Capacity) | Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for vCPU, RAM, and Storage. |
| Configuration Maximums | This widget is used to find out the VCF limits and your consumption against those limits. It displays the total number of management workload domains, VI workload domains, ESXi Hosts, vCenter Servers, and NSX-T management clusters per domain manager. |
| Topology | Gives a graphical representation of objects related to VCF. Click on each object to have an expanded view of the object details. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Domain Summary | Displays the domain name, type, cluster, virtual machine, datastore details, tanzu k8s cluster, namespace, and pods. |

## VCF Domain Summary Tab

The VCF Domain Summary tab provides details on the domain's overall inventory of the organization including key components, domain health, topology, alerts, and so on.

Where To View VCF Domain Summary Tab

From the left menu, click **Data Sources › Integrations › Accounts**. Click VMware Cloud Foundation, click the vertical ellipsis against the VCF account, and then select **Object Details**. Click the **Domain** object type and select any one of the domains in the list.

> **NOTE**
> The VCF Domain Summary tab displays the total number of domains which includes both configured and unconfigured domains.

**Table 334: Domain Summary Options**

| Option | Description |
|---|---|
| Object Summary | Displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | Displays a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts usually require immediate attention.<br>• Warning alerts indicate that you must look into any problems shortly.<br>• Info alerts indicate that you can reclaim resources. |
| Consumer | Displays the number of active VMs for the selected object. You can also view the usage details for the virtual machine, vCPU, RAM, and Provision. |
| Provider (Usable Capacity) | Displays the details of available resources for the selected object. You can view the number of hosts and capacity remaining for vCPU, RAM, Storage. |
| Topology | Displays a graphical representation of objects related to the domains. Click on each object to have an expanded view of the object details. |
| Distributed Switch | Dsiplays an overview of the state of the distributed switches. It displays the distributed switch name, version, total number of hosts, the maximum number of ports, and number of used ports. |
| Cluster Summary | Displays an overview of the state of the existing clusters. It displays the cluster name, ESXi host, virtual machine, capacity remaining, time remaining, and VM remaining. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Datastore Summary | Displays an overview of the state of the existing datastores. It displays the datastore name, capacity, virtual machine, capacity remaining, and time remaining. |

## VCF World Summary Tab

The VCF World Summary tab provides details on the VCF World inventory of the organization including key components, recommended actions for the VCF organization, and the health of the object types at the world level.

Where To View the World Summary Tab

From the left menu, click **Data Sources › Integrations › Accounts**. Click VMware Cloud Foundation, click the vertical ellipsis against the VFC account, and then select **Object Details**. Select the **VCF World** object.

**Table 335: World Summary Options**

| Option | Description |
|---|---|
| Recommended Actions | Displays details such as object type, number of accounts, and the recommended actions for each account.<br>This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts usually require immediate attention.<br>• Warning alerts indicate that you must look into any problems shortly.<br>• Normal and Unknown alerts indicate alerts you can look at later. |
| World | Displays the health badge for the VCF World object type. You can view the average health of all group members, the compliance, efficiency, risk, workload, and the total count of critical, immediate, and normal compliance. |
| Active Alerts | Displays the visual indicator of the alert status for the following alert types. You can view the name of the object and alert description, alert type, and subtype. The time when the alert was generated, the suggested fix, and when the action was taken upon the alert. |
| Information | Displays the account information of the selected object. |

**Summary Tabs for VMware Cloud on AWS**

The summary tab of VMC provides account details, overall SDDC inventory of the the organization and the key components, aggregated metrics, alerts, and so on.

## VMC Summary Tab

The VMC Summary tab provides organization details, overall SDDC inventory of the organization including key components, bill summary, and so on.

Where To View VMC Summary

From the left menu, click **Administration › Integrations › Accounts**. Click **VMware Cloud on AWS**, click the vertical ellipsis against the **VMware Cloud on AWS** account, and then select **Object Details**.

**Table 336: VMC Summary Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts usually require immediate attention.<br>• Warning alerts indicate that you must look into any problems shortly.<br>• Info alerts indicate that you can reclaim resources. |
| Consumer | Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, CPU, and Memory. |
| Provider (Usable Capacity) | Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage. |
| Bill Summary | Displays the bill name, total expense, outstanding expense, commit expense, and demand expense. |
| Configuration Maximums | This widget is used to find out the VMC limits and your consumption against those limits. It displays details of ESXi maximums, Elastic IP Addresses, and SDDCs per organization. |
| Topology | Gives a graphical representation of objects related to VMC. Click on each object to have an expanded view of the object details. |
| SDDC Summary | Displays the SDDC name, cluster, ESXi host, virtual machine, and datastore details. Click the SDDC name to view VMC SDDC Summary Tab. |

## VMC SDDC Summary Tab

The VMC SDDC Summary tab provides details on the overall SDDC inventory of the organization including key components, SDDC health, maximums, alerts, and so on.

Where To View VMC-D Organization Summary Tab

From the left menu, click **Administration › Integrations › Accounts**. Click **VMware Cloud on AWS**, click the vertical ellipsis against the **VMware Cloud on AWS** account, and then select **Object Details**. Click the **SDDC** object type and select any one of the SDDCs in the list.

**NOTE**
The VMC SDDC Summary tab displays the total number of SDDCs which includes both configured and unconfigured SDDCs.

**Table 337: VMC SDDC Summary Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts usually require immediate attention.<br>• Warning alerts indicate that you must look into any problems shortly.<br>• Info alerts indicate that you can reclaim resources. |
| Consumer | Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, CPU, and Memory. |
| Provider (Usable Capacity) | Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage. |
| Configuration Maximums | This widget is used to find out the SDDC limits and your consumption against those limits. It displays details of VPC, Cluster, ESXi, and Virtual Machine maximums. |
| Topology | Gives a graphical representation of objects related to SDDC. Click on each object to have an expanded view of the object details. |
| Cluster Summary | This widget provides an overview of the state of the existing clusters. It displays the cluster name, ESXi host, virtual machine, capacity remaining, time remaining, and VM remaining. |
| Datastore Summary | This widget provides an overview of the state of the existing datastores. It displays the datastore name, capacity, virtual machine, capacity remaining, and time remaining. |

## VMC World Summary Tab

The VMC World Summary tab provides details on the VMC World inventory of the organization including key components, recommended actions for the VMC organization, and the health of the object types at the world level.

Where To View the World Summary Tab

From the left menu, click **Administration › Integrations › Accounts**. Click VMware Cloud on AWS, click the vertical ellipsis against the VMC account, and then select **Object Details**. Select the **VMC World** object.

**Table 338: World Summary Options**

| Option | Description |
|---|---|
| Recommended Actions | This widget displays details such as object type, number of accounts, and the recommended actions for each account. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts that usually require immediate attention.<br>• Warning alerts indicate that you must look into any problems shortly.<br>• Normal and Unknown alerts indicate alerts you can look at later. |
| VMC World | This widget displays the health badge for the VMC World object type. You can view the average health of all group members, the compliance, efficiency, risk, workload, and the total count of critical, immediate, and normal compliance. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types. You can view the name of the object and alert description, alert type, and subtype. The time when the alert was generated, the suggested fix, and when the action was taken upon the alert. |
| Relationship | Displays the object relationship of the selected object. |

**Summary Tabs for Google Cloud VMware Engine**

The summary tabs for Google Cloud VMware Engine provide account details, overall Private Cloud inventory of the organization and the key components, aggregated metrics, alerts, and so on.

## GCVE Private Cloud Summary Tab

The Google Cloud VMware Engine Private Cloud Summary tab provides details about the overall Private Cloud inventory of the organization including key components, Private Cloud health, configuration maximums, alerts, and so on.

Where To View GCVE Private Cloud Summary Tab

From the left menu, click **Administration › Integrations › Accounts**. Click the drop-down against Google Cloud VMware Engine, click the vertical ellipsis against the **GCVE** account, and then select **Object Details**. Click the **Private Cloud** object type and select any one of the Private Clouds in the list. In this Private Cloud list, both configured and unconfigured Private Clouds for monitoring are shown.

**Table 339: GCVE Private Cloud Summary Options**

| Option | Description |
|---|---|
| Object Summary | This widget displays the details of the selected object for the specific project. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts usually require immediate attention. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Warning alerts indicate that you must look into any problems shortly.<br>• Info alerts indicate that you can reclaim resources. |
| Consumer | Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, CPU, and Memory. |
| Provider (Usable Capacity) | Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage. |
| Configuration Maximums | This widget displays configuration maximums at the Private Cloud level. |
| Topology | Gives a graphical representation of objects related to the Private Cloud. Click on each object to have an expanded view of the object details. |
| vSphere distributed switch | Displays the details of the vSphere distributed switch. |
| Cluster Summary | This widget provides an overview of the state of the existing clusters. It displays the cluster name, ESXi host, virtual machine, capacity remaining, time remaining, and VM remaining. |
| Datastore Summary | This widget provides an overview of the state of the existing datastores. It displays the data store name, capacity, virtual machine, capacity remaining, and time remaining. |

## GCVE Adapter Instance Summary Tab

The Google Cloud VMware Engine Adapter Instance Summary tab provides adapter instance details and the overall Private Cloud inventory of the adapter instance which includes key components, alerts, and so on.

Where To View GCVE Adapter Instance Summary Tab

From the left menu, click **Administration › Integrations › Accounts**. Click the drop-down against Google Cloud VMware Engine, click the vertical ellipsis against the GCVE account, and then select **Object Details**. Click the **GCVE Adapter Instance** object type and select the GCVE adapter instance from the list.

**Table 340: GCVE Adapter Instance Summary Options**

| Option | Description |
|---|---|
| Object Summary | This widget displays the details of the selected object for the specific project. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts usually require immediate attention.<br>• Warning alerts indicate that you must look into any problems shortly.<br>• Info alerts indicate that you can reclaim resources. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Consumer | Gives the number of active VMs for the selected object. You can also view the usage details for the virtual machine, vCPU, RAM, and Provision. |
| Provider (Usable Capacity) | Gives the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, and Storage. |
| Topology | Displays a graphical representation of objects related to GCVE. Click on each object to have an expanded view of the object details. |
| Private Cloud Summary | Displays the Private Cloud name, cluster, ESXi host, virtual machine, and datastore details. Click the Private Cloud name to view . |

## GCVE World Summary Tab

The Google Cloud VMware Engine World Summary tab provides details about the Google Cloud VMware Engine World inventory of the adapter instance including key components, recommended actions for the Google Cloud VMware Engine adapter instance, and the health of the object types at the World level.

Where To View the GCVE World Summary Tab

From the left menu, click **Administration › Integrations › Accounts**. Click the drop-down against Google Cloud VMware Engine, click the vertical ellipsis against the GCVE account, and then select **Object Details**. Click the **GCVE World** object type and select GCVE World from the list.

**Table 341: GCVE World Summary Options**

| Option | Description |
|---|---|
| Recommended Actions | This widget displays details such as object type, number of accounts, and the recommended actions for each account. This widget provides a visual indicator of the alert status for the following alert types.<br>• Critical and Immediate alerts usually require immediate attention.<br>• Warning alerts indicate that you must look into any problems shortly.<br>• Normal and Unknown alerts indicate alerts you can look at later. |
| GCVE World | This widget displays the health badge for the Google Cloud VMware Engine World object type. You can view the average health of all group members, the compliance, efficiency, risk, workload, and the total count of critical, immediate, and normal compliance. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types. You can view the name of the object and alert description, alert type, and subtype. The time when the alert was generated, the suggested fix, and when the action was taken upon the alert. |

### Azure VMware Solution Summary Tab

The Azure VMware Solution Summary tab provides organization details, overall Private Cloud inventory of the organization including key components, aggregated metrics, and so on.

Where To View the AVS Summary Tab

From the left menu, click **Administration › Integrations**. Click the vertical ellipses against the Azure VMware Solution Cloud Account and then select **Object Details**.

**Table 342: Azure VMware Solution Summary Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources. |
| Consumer | Provides the number of active VMs for the selected object. You can also view the details for the virtual machine, CPU, and Memory. |
| Provider (Usable Capacity) | Provides the details of available resources for the selected object. You can view the number of hosts and capacity remaining for CPU, RAM, Storage. |
| Topology | Provides a graphical representation of objects related to Azure VMware Solution. Click on each object to have an expanded view of the object details. |
| Private Cloud Summary | Displays the Private Cloud name, cluster, ESXi host, virtual machine, and datastore details. Click the Private Cloud name to view VMC SDDC Summary Tab. |

### Summary Tabs for vSAN

The summary tabs for vSAN provide account details, an overview of the state of the selected vSAN object, and the key components like health, capacity, alerts, and so on.

## vSAN Cluster Summary Tab

The vSAN Cluster tab provides an overview of the state of the selected vSAN cluster. For the selected object, the vSAN cluster tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN cluster and use that information to begin troubleshooting problems.

Where To View vSAN Cluster Summary Page

From the left menu, click **Inventory › Inventory Panel (Detailed View) › Integrations › All Objects › vSAN Adapter › vSAN Cluster**, and then select a specific **vSAN Cluster** object type.

**NOTE**
The vSAN Cluster Summary tab displayed depends on the vSAN cluster type.

vSAN OSA HCI Cluster Summary Tab

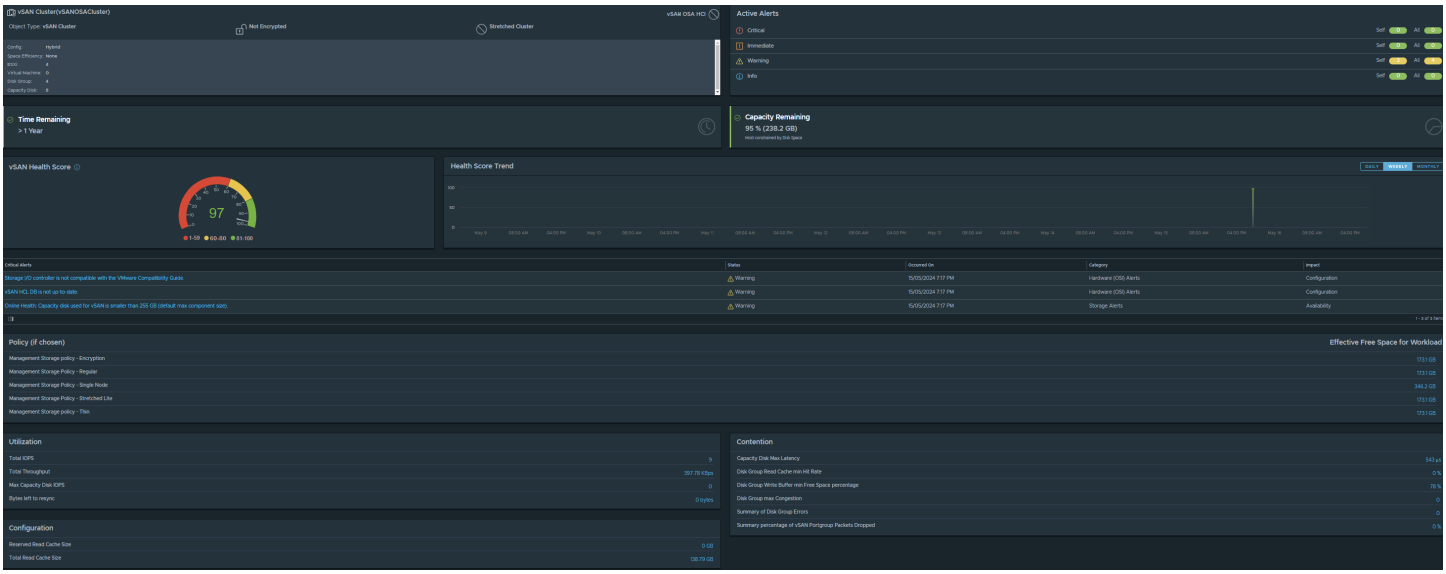


**Table 343: vSAN OSA HCI Cluster Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Critical Alerts | This widget displays alerts associated with the health score of the vSAN cluster. |
| vSAN Health Score | This widget displays the health score of the vSAN cluster. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Health Score Trend | This widget is used to find the trends in the health score of the vSAN cluster. You can view daily, weekly, and monthly health score trends. |
| Effective Free Space for Workload based on Chosen Policy | This widget displays the estimated capacity for a future workload, that is based on the selected storage policy. |
| Health Score Trend | This widget is used to find the trends in the health score of the vSAN cluster. You can view daily, weekly, and monthly health score trends. |
| Utilization | This widget is used to find out the trends in capacity used by a selected vSAN cluster as against the total capacity available. |
| Configuration | This widget displays the configuration details of the cluster. |
| Contention | This widget displays the memory contention details of the vSAN cluster. |

vSAN ESA HCI Summary Tab

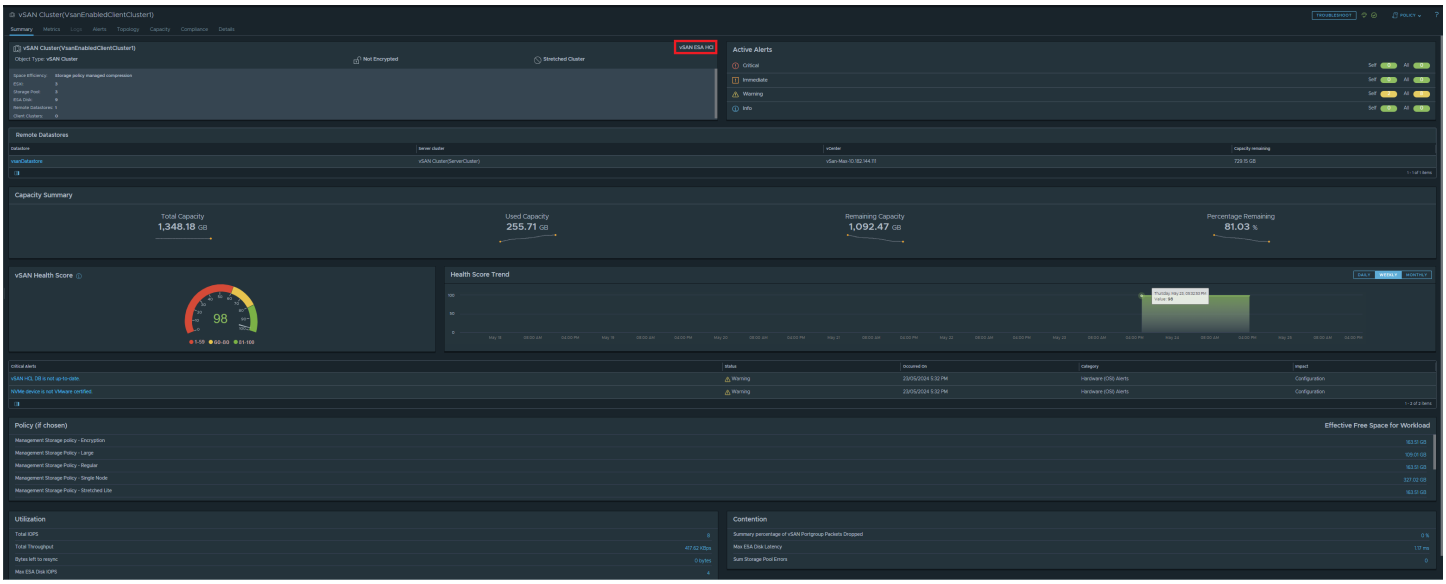


**Table 344: vSAN ESA HCI Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Effective Free Space for Workload based on Chosen Policy | This widget displays the estimated capacity for a future workload, that is based on the selected storage policy. |
| vSAN Health Score | This widget displays the health score of the vSAN cluster. |
| Health Score Trend | This widget is used to find the trends in the health score of the vSAN cluster. You can view daily, weekly, and monthly health score trends. |
| Critical Alerts | This widget displays alerts associated with the health score of the vSAN cluster. |
| Utilization | This widget is used to find out the trends in capacity used by a selected vSAN cluster as against the total capacity available. |
| Contention | This widget displays the memory contention details of the vSAN cluster. |

vSAN Compute Cluster Summary Tab

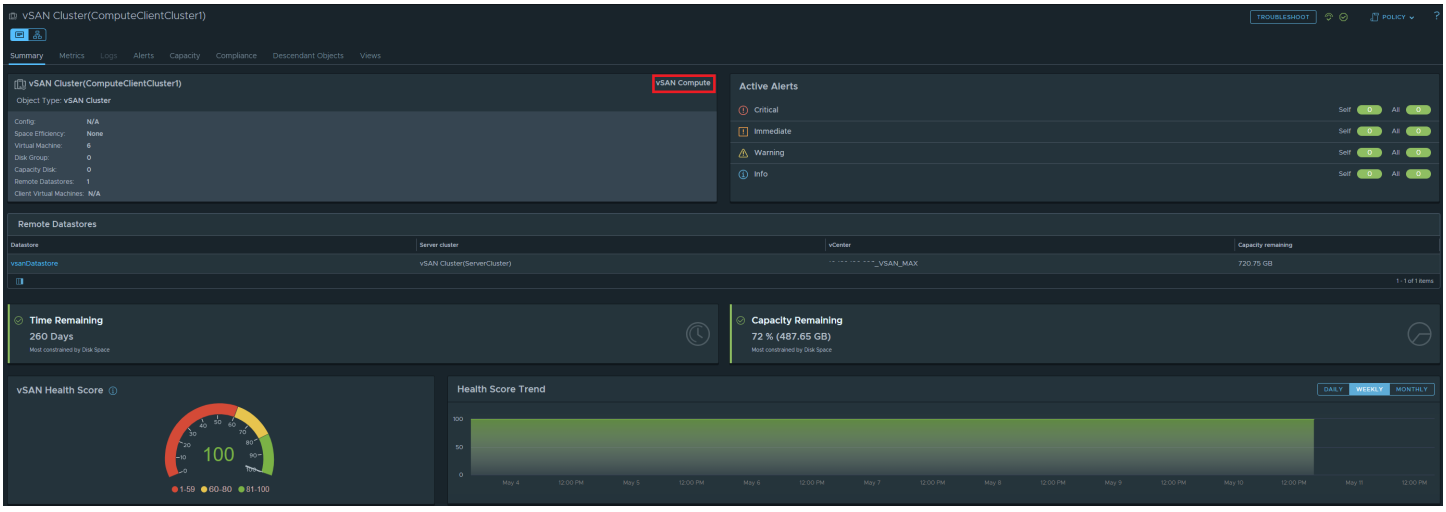

**Table 345: vSAN Compute Cluster Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Remote Datastores | Displays the remote datastores used by vSAN compute.<br><br>**NOTE**<br>vSAN compute cluster remotely mounts the datastore of vSAN Max. |
| vSAN Health Score | This widget displays the health score of the vSAN cluster. |
| Health Score Trend | This widget is used to find the trends in the health score of the vSAN cluster. You can view daily, weekly, and monthly health score trends. |

vSAN ESA Max Cluster Summary Tab

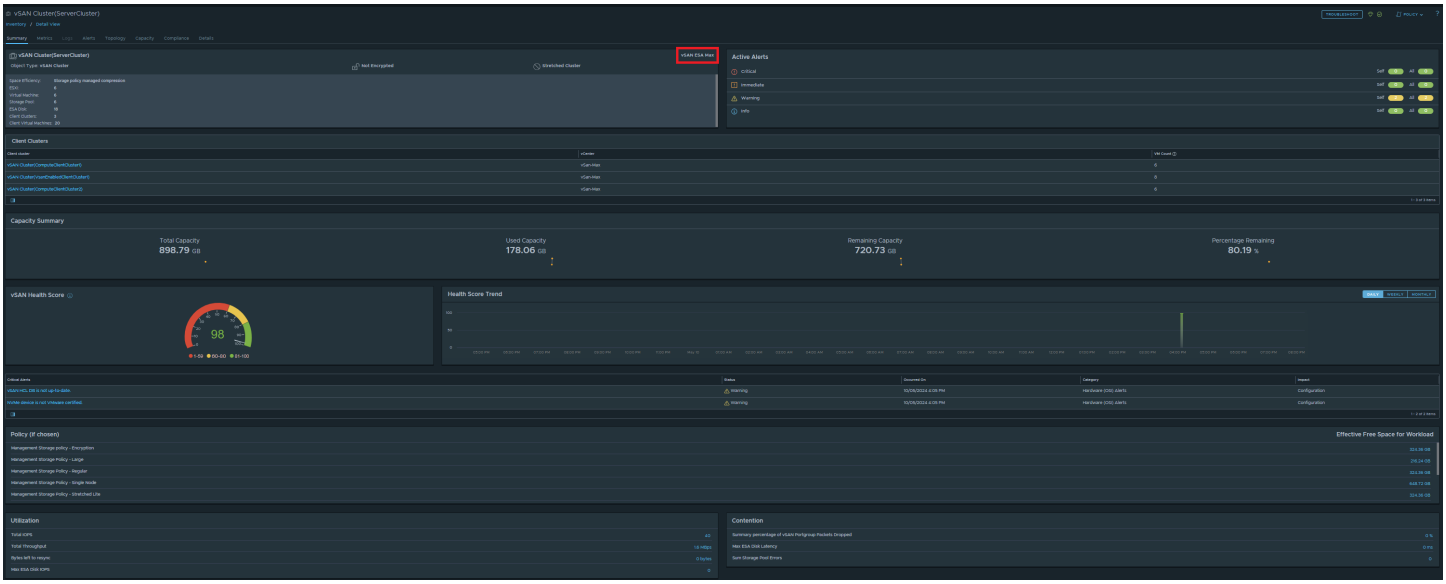

**Table 346: vSAN Max ESA Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Client Clusters | This widget displays total number of VMs using the vSAN Max cluster for each client cluster. |
| Capacity Summary | This widget displays the total capacity, used capacity, remaining capacity of your virtual environment in gigabytes. It also displays the remaining capacity in percentage. Use this widget to monitor your cluster's storage capacity to accommodate new virtual machines. |
| vSAN Health Score | This widget displays the health score of the vSAN cluster. |
| Health Score Trend | This widget is used to find the trends in the health score of the vSAN cluster. You can view daily, weekly, and monthly health score trends. |
| Critical Alerts | This widget displays alerts associated with the health score of the vSAN cluster. |
| Effective Free Space for Workload based on Chosen Policy | This widget displays the estimated capacity for a future workload, that is based on the selected storage policy. |
| Utilization | This widget is used to find out the trends in capacity used by a selected vSAN cluster as against the total capacity available. |
| Contention | This widget displays the memory contention details of the vSAN cluster. |

## vSAN Disk Group Summary Tab

The vSAN Disk Group Summary tab provides an overview of the state of the selected vSAN Disk Group. For the selected object, the vSAN Disk Group tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN Disk Group and use that information to begin troubleshooting problems.

Where To View vSAN Cluster Disk Group Summary

From the left menu, click **Inventory › Inventory Panel (Detailed View) › Integrations › All Objects › vSAN Adapter › vSAN Disk Group**, and then select a **Disk Group** object type.

**Table 347: vSAN Cluster Disk Group Summary Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Utilization | This widget is used to find out the trends in capacity used by a selected vSAN cluster disk group as against the total capacity available. |
| Contention | This widget displays the memory contention details of the vSAN cluster. |
| Resync | This widget displays the throughput and latency details for the vSAN cluster disk group. |

## vSAN Capacity Disk Summary Tab

The vSAN Capacity Disk tab provides an overview of the state of the selected vSAN capacity disk. For the selected object, the vSAN capacity disk tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN capacity disk and use that information to begin troubleshooting problems.

**Table 348: vSAN Capacity Disk Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Utilization | This widget is used to find out the trends in capacity used by a selected capacity disk as against the total capacity available. |
| Contention | This widget displays the memory contention details for the selected capacity disk. |

## vSAN Cache Disk Summary Tab

The vSAN Cache Disk tab provides an overview of the state of the selected vSAN cache disk. For the selected object, the vSAN cache disk tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN cache disk and use that information to begin troubleshooting problems.

**Table 349: vSAN Cache Disk Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Utilization | This widget is used to find out the trends in capacity used by a selected vSAN cache disk as against the total capacity available. |
| Contention | This widget displays the memory contention details for the selected cache disk. |

## vSAN Cluster Fault Domain Summary Tab

The vSAN cluster fault domain summary tab provides details about CPU, CPU Cores, Memory, Disc Space and Alerts associated with the fault domain of the vSAN cluster.

Where To View vSAN Cluster Fault Domain Summary

On the menu, click **Inventory › Inventory Panel (Detailed View) › Integrations › All Objects › vSAN Adapter › vSAN and Storage Devices › vSAN Cluster › Fault Domain**.

You can also view relationship details and heat map details for the selected vSAN fault domain. The relationship section provides information about the relationship between the objects in your vSAN cluster. The heat map helps you to identify potential problems for the objects in your vSAN fault domain.

## vSAN World Summary Tab

The vSAN World summary tab provides an overview of the state of all the vSAN clusters in your environment. For the selected object, the vSAN World Summary tab displays details about capacity remaining, cluster health score, vCenter, Datacenter, ESXi Host, vSAN Capacity, Space Efficiency, alerts, and metrics as they affect the health, risk, or efficiency of the cluster. You can use this tab to identify the vSAN ESA (Express Storage Architecture) clusters or non vSAN ESA clusters in your environment. You can evaluate the impact of alerts on vSAN clusters and use that information to begin troubleshooting problems.

Where To View vSAN World Summary Page

From the left menu, click **Inventory › Inventory Panel (Detailed View) › Integrations › All Objects › vSAN Adapter › vSAN World**, and then select a **vSAN World** object type.

**Table 350: vSAN World Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object, you can also view the number of resources associated with the selected object. It displays the total number of vSAN clusters and the number of each vSAN cluster type. It also displays the total number of ESXi hosts, virtual machines, and vSAN Datastore at the vSAN World level. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| vSAN Cluster Summary | vSAN World covers all the vSAN clusters in the environment. The cluster information is provided in a tabular format with the following columns:<br>• vSAN Cluster: Name of the vSAN cluster.<br>• Cluster Health Score: Displays the overall heath score of the cluster.<br>• vCenter: Displays the vCenter Server associated with the cluster.<br>• Datacenter: Displays the datacenter used by the cluster.<br>• vSAN Type: Displays the vSAN cluster type. For more information, see vSAN Cluster Summary Tab.<br>• ESXi Host: Displays the number of ESXi hosts used in each cluster.<br>• vSAN Capacity: Displays the total capacity in gigabytes.<br>• Capacity Remaining: Displays the remaining capacity in each cluster in gigabytes.<br>• Space Efficiency: Displays the space efficiency feature the cluster has subscribed to during configuration.<br>• Stretched Cluster: Displays if the stretched cluster option is activated or deactivated. |

## vSAN Adapter Summary Tab

The vSAN Adapter summary tab provides an overview of the state of all the vSAN adapters in your environment. For the selected object, the vSAN adapter Summary tab displays details about capacity remaining, cluster health score, vCenter, Datacenter, ESXi Host, vSAN Capacity, Space Efficiency, alerts, and metrics as they affect the health, risk, or efficiency of the adapter. You can evaluate the impact of alerts on vSAN adapters and use that information to begin troubleshooting problems.

Where To View vSAN Adapter Summary Page

From the left menu, click **Inventory › Inventory Panel (Detailed View) › Integrations › All Objects › vSAN Adapter › vSAN Adapter Instance**, and then select the **Adapter Instance** object type.

**Table 351: vSAN Adapter Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object, you can also view the number of resources associated with the selected object. The vSAN ESA (Express Storage Architecture) column indicates whether the cluster is ESA activated or not. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| vSAN Cluster Summary | Each vSAN adapter instance can have one or more clusters. The cluster information is provided in a tabular format with the following columns:<br>• vSAN Cluster: Name of the vSAN cluster.<br>• Cluster Health Score: Displays the overall heath score of the cluster.<br>• Datacenter: Displays the datacenter used by the cluster.<br>• vSAN Type: Displays the vSAN cluster type. For more information, see vSAN Cluster Summary Tab.<br>• ESXi Host: Displays the number of ESXi hosts used in each cluster.<br>• vSAN Capacity: Displays the total capacity in gigabytes.<br>• Capacity Remaining: Displays the remaining capacity in each cluster in gigabytes.<br>• Space Efficiency: Displays the space efficiency feature the cluster has subscribed to during configuration.<br>• Stretched Cluster: Displays if the stretched cluster option is activated or deactivated. |

## vSAN Storage Pool Summary Tab

The vSAN Storage Pool summary tab provides an overview of the state of all the vSAN Storage pool in your environment. For the selected object, the vSAN Storage Pool Summary tab displays details about capacity remaining, vCenter, Datacenter, ESXi Host, vSAN Capacity, Space Efficiency, alerts, and metrics as they affect the health, risk, or efficiency of the cluster.You can evaluate the impact of alerts on the Storage Pool and use that information to begin troubleshooting problems.

Where To View vSAN Storage Pool Summary Page

From the left menu, click **Inventory › Inventory Panel (Detailed View) › Integrations › All Objects › vSAN Adapter › vSAN Storage Pool**, and then select a **Storage Pool** object type.

**Table 352: vSAN Storage Pool Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Object Summary | This widget displays the details of the selected object. The widget also displays the number of resources associated with the selected object. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources.<br><br>To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Contention | This widget displays the memory contention details of the vSAN Storage Pool. |

## vSAN ESA Disk Summary Tab

The vSAN ESA Disk Summary tab provides an overview of the state of the selected vSAN ESA cluster. For the selected object, the vSAN ESA cluster tab displays the alerts, time remaining, capacity remaining, utilization, configuration, and metrics as they affect the health, risk, or efficiency. You can use this tab to evaluate the impact that alerts are having on the vSAN ESA cluster and use that information to begin troubleshooting problems.

Where To View vSAN ESA Disk Summary Page

From the left menu, click **Inventory › Inventory Panel (Detailed View) › Integrations › All Objects › vSAN Adapter › vSAN ESA Disk**, and then click a **vSAN ESA Disk** object type.

**Table 353: vSAN ESA Disk Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot | Start the Troubleshooting Workbench with the current object in context. |
| Object Summary | This widget displays the details of the selected object. The widget also displays the disk details like, size, vendor, model, queue depth, and disk type. |
| Active Alerts | This widget provides a visual indicator of the alert status for the following alert types.<br>• Health alerts that usually require immediate attention.<br>• Risk alerts indicating that you must look into any problems shortly.<br>• Efficiency alerts indicating that you can reclaim resources. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | To see the alerts for the object, click the labels of the alert. |
| Time Remaining | This widget displays the number of days remaining till the projected resource utilization crosses the threshold for the usable capacity. |
| Capacity Remaining | This widget displays the unused capacity of your virtual environment to accommodate new virtual machines. |
| Utilization | This widget is used to find out the trends in capacity used by a selected vSAN cluster as against the total capacity available. |
| Contention | This widget displays the memory contention details of the vSAN cluster. |

**Business Application Summary Tab**

The Business Application Summary tab provides an overview of the selected Business Application. Each card on the page provides additional context such as overview, active alerts, KPIs, and object relationships between members of the Business Application.

Understanding the Business Application Summary Tab

**Table 354: Business Application Summary Tab Options**

| Option | Description |
|---|---|
| Troubleshoot button | Launch the Troubleshooting Workbench with the current Business Application in context. |
| Object Summary card | Displays the details of the current Business Application.<br>• Source<br>• Application Tag<br>• Environment<br>• Business Criticality<br>• Description |
| Business Application KPIs card | Displays the availability of the Business Application calculated per calendar month. |
| Active Alerts card | Displays a summary of active alerts divided by **Self** and **All**, which are triggered on the Business Application and the child objects.<br>For more details about the alerts by object, see the relationship diagram, or the relationship table section. |
| Relationship view diagram | This section provides two types of relationship diagram views:<br>• Application View. This is the same relationship diagram which you can view in the Business Application page. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Infrastructure View. Displays a graph of the Virtual Machines in the Business Application and their vSphere infrastructure.<br><br>You can switch between the horizontal view and vertical view by toggling the orientation icon on the right of the relationship view diagram.<br><br>Move your cursor over an object in the relationship diagram to view more details in a transient pop-up dialog box which appears. |
| Relationship view table | Displays the following details for the objects listed in the Business Application relationship diagram:<br>• Name<br>• Object Type<br>• Capacity Remaining<br>• Alerts |

**Evaluating Metric Information**

The **Metrics** tab provides a relationship map and user-defined metric charts. The topological map helps you evaluate objects in the context of their place in your environment topology. The metric charts are based on the metrics for the selected object that you think helps identify the possible cause of a problem in your environment.

Although you might be investigating problems with a single object, for example, a host system, the relationship map allows you to see the host in the context of parent and child objects. It also works as a hierarchical navigation system. If you double-click an object in the map, that object becomes the focus of the map. The available metrics for the object become active in the lower-left pane.

> **NOTE**
> The yellow diamond icon next to the metric indicates dynamic threshold breach, the blue diamond icon next to the metric indicates the metric value is within threshold.

You can also build your own set of metric charts. You select the objects and metrics that provide you with a detailed view of changes to different metrics for a single object, or for related objects over time.

Where available, the **Metrics** tab provides pre-defined sets of metrics to help you when looking at a specific aspect of an object. For example, if you have a problem with a host, access the most relevant information about the host by looking at the metrics displayed in the pre-defined lists. You can edit these groups of metrics, and create additional groups, by dragging and dropping metrics and properties from the All Metrics and All Properties lists.

**Where You Find the Metrics Tab**

- From the left menu, click **Inventory**, then select a group, custom data center, application, or inventory object.

**Create Metric Charts When You Troubleshoot a Virtual Machine Problem**

You create a custom group of metric charts when you troubleshoot a problem with a virtual machine so that you can compare different metrics. The level of detail that you can create using the **Metrics** tab, can contribute significantly to your effort to find the root cause of a problem.

As an administrator investigating a performance problem with a virtual machine, you determined that you must see detailed charts about the following reported symptoms.

- Guest file system overall disk space usage reaching critical limit
- Guest partition disk space usage

The following method of evaluating problems using the **Metrics** tab is provided as an example for using VMware Aria OperationsVMware Cloud Foundation Operations and is not definitive. Your troubleshooting skills and your knowledge of the particulars of your environment determine which methods work for you.

1. Enter the name of the virtual machine in the **Search** text box on the menu bar.
   In this example, the virtual machine name is `sales-10-dk`.
2. Click the **Metrics** tab.
3. In the relationship topology map, click the virtual machine, **dk-new-10**.
   The metrics list, located in the left of the center pane, displays virtual machine metrics.
4. On the chart toolbar, click **Date Control** and select a time that is on or before the symptoms were triggered.
5. Add metric charts to the display area for the virtual machine.
   a) In the metric list, select **Guest Files System Stats › Total Guest File System Free (GB)** and double-click the metric name.
   b) To add the guest partition, for example, C:\, select **Guest Files System Stats › C:\ › Guest File System Free (GB)** and double-click the metric name.
   c) To add disk space for comparison, select **Disk Space › Capacity Remaining (%)** and double-click the metric name.
6. Compare the charts.

   You can see a decrease in the file system free space, and that the virtual machine disk space capacity remaining is decreasing at a steady rate. You determine that you must add disk space to the virtual machine. However, you do not know if the datastore can support the change to the virtual machine.

7. Add the datastore capacity chart to the charts.
   a) In the topology map, double-click the host.
      The topology map refreshes with the host as the focus object.
   b) Click the datastore.
   c) In the metric list, which is updated to display datastore metrics, select **Capacity › Available Space (GB)** and double-click the metric name.
8. To determine if sufficient capacity is available on the datastore to support increasing the disk space on the virtual machine, review the datastore capacity chart.

You know that you must increase the size of the virtual disk on the virtual machine.

Expand the virtual disk on the virtual machine and assign it to stressed partitions. Click **Actions**, on the object title bar, and view the virtual machine in the vSphere Web Client.

### Troubleshooting with the Metrics Tab

The **Metrics** tab provides a relationship graph and metric charts. The relationship graph helps you evaluate objects in the context of their place in your environment topology. Metric charts are based on the metrics for the active map object that you think can help you identify the cause of a problem.

## How the Metrics Tab Works

You can double-click any object in the graph and view the specific parent-child objects for the focus object. If you point to an object icon, you can see the health, risk, and efficiency details. You can also click the **Alerts** link for the number of generated alerts. Click the purple icon to view the child relationships of the object. If you double-click an object icon, the selected object becomes the focus of the map. The graph is updated for the selected object, and the metrics list shows only the metrics for the selected object.

Using the metrics list, you create charts based on metrics that you think can help you investigate problems. You customize the charts to evaluate the data in detail. To save the configured charts, you create a dashboard using the toolbar option.

Where available, the metrics list also displays pre-defined groups of metrics that contain the most relevant metrics for the selected object. You can edit these groups, and create your own customized groups of metrics by dragging and dropping metrics and properties from the All Metrics and All Properties lists.

## Where You Find the Metrics Tab

- From the left menu, click **Inventory**, then select a group, custom data center, application, or inventory object.
- Alternatively, click **Inventory**, then use the hierarchies in the left pane to locate the objects that you want.

## Edit Metric State

The edit metric state option allows you activate, deactivate, or inherit metric state, DT state and KPI state. You can use the All Filter option to search and edit the metrics based on the following criteria.

- Name
- State
- Instanced State
- DT State
- KPI State
- Show Only Supermetrics
- Local Changes
- Unsaved Changes

## Metrics Options

The options include the graph toolbar, the metric selector options, the metric charts toolbar, and the toolbar on each chart.

**Table 355: Relationship Map Toolbar Options**

| Option | Description |
|---|---|
| Back to initial object | Returns the map to original object if you double-clicked on an icon to examine another object. |
| Vertical/Horizontal | Displays a vertical or horizontal view of the graph or tree view. |
| Hide Text/Show Text | Hides or displays the object names. |
| Standard View/Fit View | The **Standard View** option fixes the view to a specific zoom level.<br>The **Fit View** option adjusts the graph or tree view to fit the screen. |

**Table 356: Specific Functions from the Relationship Map**

| | |
|---|---|
| Hide Node | Hover your cursor over an object and click the **Hide Node** icon to hide the specific node. |
| Show Peers/Hide Peers | Hover your cursor over the selected object in the topology pane and select **ShowPeers/Hide Peers** to view or hide other objects of the same type that exist under the parent object. |
| Pagination/Next Page/Previous Page | At the bottom of the topological view you can view the page number and also move to the next/previous page. |
| Filter | At the bottom of the topological view you can use the filter option to search for objects. |
| Arrows | Use the arrows on each object to view relationships of each object. |

The chart options are used to limit the metric list.

**Table 357: Metric Chart Selector**

| Option | Description |
|---|---|
| Show collecting metrics | Updates the list to display only the currently collected metrics for the object. |
| Show previewable super metrics | Updates the list to display super metrics for the object.<br>**NOTE**<br>The super metrics only appear if the super metric is associated with the object, see Create a Super Metric topic in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* . |
| Actions | Click the **Actions** icon to configure metric groups. Verify that you hold the PowerUser or administrator role.<br>• **Add Group**. To add metrics or properties to the group, expand any of the metric groups, and drag one or more metrics to the group.<br>• **Remove Group(s)**. To remove one or more groups.<br>• **Rename Group**. To enter a new name for the group.<br>• **Remove Metric(s) from Group(s)**. To remove one or more metrics or properties from one or more groups, hold down the Ctrl key, and select the metrics or properties that you want to remove. |
| Search | Use a word search to limit the number of items that appear in the list. |
| Time Range | Filters the metrics to show only the ones that have received data in the selected time range. |
| Metric list | Double-click a metric to populate the chart window. To populate the chart window with a separate chart for each of the metrics in the group, double-click a metric group. |

To visualize the specific metric data over time, and compare the results for different metrics, select different combinations of options.

**Table 358: Metric Chart Toolbar**

| Option | Description |
|---|---|
| Split Charts | Displays each metric in a separate chart. |
| Stacked Chart | Consolidates all charts into one chart. This chart is useful for seeing how the total or sum of the metric values vary over time. To view the stacked chart, ensure that the split chart option is turned off. |
| Y Axis | Shows or hides the Y-axis scale. |
| Metric Chart | Shows or hides the line that connects the data points on the chart. |
| Trend Line | Shows or hides the line and data points that represents the metric trend. The trend line filters out metric noise along the timeline by plotting each data point relative to the average of its adjoining data points. |
| Dynamic Thresholds | Shows or hides the calculated dynamic threshold values for a 24-hour period. |
| Show Entire Period Dynamic Thresholds | Shows or hides dynamic thresholds for the entire time period of the graph. |
| Anomalies | Shows or hides anomalies. Time periods when the metric violates a threshold are shaded. Anomalies are generated when a metric crosses a dynamic or static threshold, either above or below. |
| Show Data Point Tips | Shows or hides the data point tooltips when you hover the mouse over a data point in the chart. |
| Zoom All Charts | Resizes all the charts that are open in the chart pane based on the area captured when you use the range selector.<br><br>You can switch between this option and **Zoom the View**. |
| Zoom the View | Resizes the current chart when you use the range selector. |
| Pan | When you are in zoom mode, allows you to drag the enlarged section of the chart so that you can view higher or lower, earlier or later values for the metric. |
| Show Data Values | Activates the data point tooltips if you switched to a zoom or pan option. **Show Data Point Tips** must be activated. |
| Refresh Charts | Reloads the charts with current data. |
| Date Controls | Opens the date selector.<br><br>Use the date selector to limit the data that appears in each chart to the time period you are examining. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Near Real-Time Monitoring | Use the near real-time monitoring option to view the data collected for one or more selected metrics. The real-time data ranges from 24 hours to three days. |
| Generate Dashboard | Saves the current charts as a dashboard. |
| Remove All | Removes all the charts from the chart pane, allowing to you begin constructing a new set of charts. |

Manage individual charts with the toolbar options.

**Table 359: Individual Metric Charts Toolbar**

| Option | Description |
|---|---|
| Navigation | If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application. |
| Correlation | Runs metric correlation using the following options:<br><br>**Self-Metrics**: Runs metric correlation on all metrics for the selected object, to find metrics of similar, or opposite behavioral change for the same time period. The instanced metrics are not assessed in the self-metrics correlation method.<br><br>**Peers**: Runs metric correlation on the same metrics for all peer objects, to find the same metrics with behavioral changes within peer objects. Peer objects are the direct child objects of the parent for the selected objects. The child objects have the same object type.<br><br>**NOTE**<br>The correlation results only appears if there are at least 11 data points and the time range is within the three months period to run the metric correlation.<br><br>**Scope**: Runs metric correlation on all metrics for the selected object with the selected scope, to find metrics of similar, or opposite behavioral change for the same time period. The instanced metrics are not assessed in the scope correlation method.<br><br>To run metrics correlation using real-time data, click the **Near Real-Time Monitoring** icon before proceeding.<br><br>After you run the correlation, the results are displayed in the **Correlation** window. By default, only the first 10 results for correlated metrics are displayed. To view the full list, click **Show More**. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | You can zoom in to view the correlated metrics and also pin them so that they appear in the preview section of the **All Metrics** tab.<br><br>**NOTE**<br>During the correlation process, some metrics are left out. For example, the badge and VMware Aria OperationsVMware Cloud Foundation Operations generated metrics. By default, the instanced metrics are omitted, except those in the **Aggregate of all instances** group. |
| Save a Snapshot | Creates a PNG file of the current chart. The image is the size that appears on your screen.<br><br>You can retrieve the file in your browser's download folder. |
| Save a Full Screen Snapshot | Downloads the current graph image as a full-page PNG file, which you can display or save.<br><br>You can retrieve the file in your browser's download folder. |
| Create an Alert Definition | Allows you to create an alert for an object type or metric in a quick and easier way. For details, see Create a Simple Alert Definition section in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* . |
| Download comma-separated data | Creates a CSV file that includes the data in the current chart.<br><br>You can retrieve the file in your browser's download folder. |
| Scales | You can choose a scale for a stacked chart.<br>• Select **Linear** to view a chart in which the Y-axis scale increases in a linear manner. For example, the Y-axis can have ranges from 0 to 100, 100 to 200, 200 to 300, and so on.<br>• Select **Logarithmic** to view a chart in which the Y-axis scale increases in a logarithmic manner. For example, the Y axis can have ranges from 10 to 20, 20 to 300, 300 to 4000, and so on. This scale gives a better visibility of minimum and maximum values in the chart when you have a large range of metric values.<br><br>**NOTE**<br>If you select a logarithmic scale, the chart does not display data points for metric values less than or equal to 0, which leads to gaps in the graph.<br><br>• Select **Combined** to view overlapping graphs for the metrics. The chart uses individual scales for each graph |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | instead of using a relative scale, and displays a combined view of the graphs.<br>• Select **Combined by Unit** to view a chart that groups the graphs for similar metric units together. The chart uses a common scale for the combined graphs. |
| Move Down | Moves the chart down one position. |
| Move Up | Moves the chart up one position. |
| Close | Deletes the chart. |
| Vertical resize | Resizes the height of a graph in the chart. |
| **Remove** icon next to each metric name in a stacked chart | Removes the graph for the metric from the chart. |

### Investigating Object Alerts

The **Alerts** tab provides a list of generated alerts for the currently selected object. When you are working with objects, reviewing and responding to generated alerts on the **Alert** tab helps you manage problems in your environment.

The alerts notify you when a problem occurs in your environment based on configured alert definitions. Object alerts are useful to you as an investigative tool in two ways. They can provide you with early notification about problems in your environment before a user calls you to report a problem. As well, object alerts can provide information about the object that you can use when troubleshooting general or reported problems.

As you review the **Alerts** tab, you can add ancestors and descendants to the list to broaden your view of the alerts. You can see if alerts on the current object affect other objects. Conversely, you can examine how problems reflected in alerts on other objects affect the current object.

Depending on the practices and workflows of your infrastructure operations team, you can use the object **Alerts** tab to manage generated alerts on individual objects.

- Take ownership of alerts so that your team knows that you are working to resolve the problem.
- Suspend an alert so that is temporarily excluded from affecting the Health, Risk, or Efficiency state of the object while you investigate the problem.
- Cancel alerts that you know are a result of a deliberate action. For example, a network card is removed from a host for replacement. Also cancel alerts that are known issues that you cannot resolve currently because of resource constraints. Canceling an alert that is generated because of only message event or metric event symptoms cancels the alert permanently. If the underlying metric or property condition remains true, canceling an alert that is generated because of metric, super metric, or property symptoms can result in the alert being regenerated . It is only effective to cancel alerts generated because of message event or metric event symptoms.

Investigating and resolving alerts helps you provide the best possible environment to your customers.

### Alerts Tab

The Alerts tab is a list of all the alerts generated for the selected object, group, or application. Use the alerts list to evaluate the number of generated alerts for the object so that you can begin resolving them.

#### How the Alerts Tab Works

All the active alerts for the selected object appear in the list. By default, the system groups the alerts by Definition. You can select multiple rows in the list using Shift+click, Control+click. Modify the filter if you want to see inactive alerts.

Manage the alerts in the list using the toolbar options. Click the **alert name** to see the alert details for the affected object. The alert details appear on the right, including the symptoms triggered with the alert. The system offers recommendations for addressing the alert and links to additional information. A **Run Action** button might appear in the details. Point to the button to learn what recommendation is performed if you click the button. To return to the list view, click the **X** at the top right of the alert details.

To see the object details, click the **Summary** Tab.

### Where You Find the Alerts Tab

- From the left menu, browse to the **Global Inventory** page, and then select a group, custom data center, application, or inventory object. Click the **Alerts** tab.
- In the menu, select **Search** and locate the object of interest. Click the **object** to display the object's **Summary** tab. Click the **Alerts** tabs.

### Alerts Display Options

The alert options include toolbar and data grid options. Use the toolbar options to sort the alert list and to cancel, suspend, or manage ownership. Additional toolbar options enable you to review parent and child alerts related to the alert you are reviewing. Use the data grid to view the alerts and alert details.

**Table 360: Actions Menu**

| Option | Description |
|---|---|
| Actions menu | Select an alert from the list to turn on the Actions menu, then select an option from the menu. |
| Menu Options: | |
| Cancel Alert | Cancels the selected alerts. If you configure the alert list to display only active alerts, the canceled alert is removed from the list.<br><br>You cancel alerts when you do not need to address them. Canceling the alert does not cancel the underlying condition that generated the alert. Canceling alerts is effective if the alert is generated by triggered fault and event symptoms because these symptoms are triggered again only when subsequent faults or events occur on the monitored objects. If the alert is generated based on metric or property symptoms, the alert is canceled only until the next collection and analysis cycle. If the violating values are still present, the alert is generated again. |
| Delete Canceled Alerts | Delete canceled (inactive) alerts by making a group selection or by individually selecting alerts. You cannot delete active alerts. |
| Suspend | Suspend an alert for a specified number of minutes.<br><br>You suspend alerts when you are investigating an alert and do not want the alert to affect the health, risk, or efficiency of the object while you are working. If the problem persists after the elapsed time, the alert is reactivated and it will again affect the health, risk, or efficiency of the object. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | The user who suspends the alert becomes the assigned owner. |
| Assign to | Assign the alert to a user. You can search for a specific username and click **Save** to assign the alert to the selected user. |
| Take Ownership | As the current user, you make yourself the owner of the alert.<br><br>You can only take ownership of an alert, you cannot assign ownership. |
| Release Ownership | Alert is released from all ownership. |
| Go to Alert Definition | Switches to the Alert Definitions page, with the definition for the previously selected alert displayed. |
| Deactivate | Offers two options for disabling the alert:<br>Deactivate the alert in all policies: this deactivates the alert for all objects for all the policies.<br><br>Deactivate Alert in Selected Policies: this deactivates the alert for objects having the selected policy. This method works only for objects with alerts. |
| Open an external application | Actions you can run on the selected object.<br>For example, Open Virtual Machine in vSphere Client. |

**Table 361: Include Menu**

| Options | Description |
|---|---|
| Self | The selected object. |
| Parents <options> | Displays the alerts for the ancestors of the selected object.<br><br>Parents in this instance include the parents, grandparents, and so on, of the object. For example, the parents of a host are a folder, storage pod, cluster, and data center instance. |
| Children <options> | Displays the alerts for the descendants of the selected object.<br><br>Children in this instance include the children and grandchildren of the object. For example, the descendants of a host are datastores, resources pools, and virtual machines. |
| Peer | Shows or hides alerts for objects with the same kind of the impacted object which also share the same parent. |

**Table 362: Group By Options**

| Option | Description |
|---|---|
| None | Alerts are not sorted into specific groupings. |
| Time | Group alerts by time triggered. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Criticality | Group alerts by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. See also Criticality in the "All Alerts Data Grid Options" table, below. |
| Definition | Group alerts by definition, that is, group like alerts together. Selected by default. |
| Object Type | Group alerts by the type of object that triggered the alert. For example, group alerts on hosts together. |
| Scope | Group alerts by scope. You can search for alerts within the selected scope. |

**Table 363: Quick Filters**

| | |
|---|---|
| Filtering options | The advanced search and filter lets you search for a symptom by:<br>• Alert ID<br>• Alert<br>• Owner<br>• Impact<br>• Alert Type<br>• Alert Subtype<br>• Status<br>• Criticality<br>• Triggered On<br>• Control State<br>• Object Type<br>• Created On<br>• Updated On<br>• Canceled On<br>• Action |

**Table 364: Alerts Data Grid**

| Option | Description |
|---|---|
| Criticality | Criticality is the level of importance of the alert in your environment. The alert criticality appears in a tooltip when you hover the mouse over the criticality icon.<br><br>The level is based on the level assigned when the alert definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**.<br><br>The possible values include:<br>• Critical<br>• Immediate |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Warning<br>• Information |
| Alert | Name of the alert definition that generated the alert.<br><br>Click the alert name to view the alert details tabs where you can begin troubleshooting the alert. |
| Created On | Date and time when the alert was generated. |
| Status | Current state of the alert.<br><br>Possible values include Active or Canceled. |
| Alert Type | Describes the type of alert that triggered on the selected object, and helps you categorize the alerts so that you can assign certain types of alerts to specific system administrators. For example, Application, Virtualization/ Hypervisor, Hardware, Storage, Network, Administrative, and Findings. |
| Alert Subtype | Describes additional information about the type of alert that triggered on the selected object, and helps you categorize the alerts to a more detailed level than Alert Type, so that you can assign certain types of alerts to specific system administrators. For example, Availability, Performance, Capacity, Compliance, and Configuration. |
| Importance | Displays the priority of the alert. The importance level of the alert is determined using a smart ranking algorithm. |

**Alert Details Tab**

| Section | Description |
|---|---|
| Recommendations | View recommendations for the alert. Click **<** or **>** to cycle through the recommendations. To resolve the alert, click the **Run Action** button if it appears. |
| Other Recommendations | Collapse the section to view additional recommendations. See the links in the **Need More Information?** section to view additional metrics, events, or other details that appear as a link. |
| Alert Basis | |
| Active Only | This option is activated by default. When activated, all active symptoms/conditions that were met for the alert are displayed. When deactivated, all the symptoms/conditions of an alert are displayed. |
| Symptoms | View the symptoms that triggered the alert. Collapse each symptom to view additional information. |
| Conditions | View the conditions that triggered the alert. Collapse each condition to view additional information. |
| Troubleshoot with Logs | Opens the **Logs** tab in the Troubleshooting Workbench, with the current object in context. Displays logs from 15 minutes before when the alert was triggered. |

*Table continued on next page*

*Continued from previous page*

| Section | Description |
|---|---|
| Notes | Enter your notes about the alert and click **Submit** to save. |
| Close | Click the X icon to close the alert details tab. |

**Related Alerts Tab**

The **Related Scope** displayed on the right, shows the objects that are one level above and one level below the object on which the alert was triggered. This topology is fixed. You cannot change the scope in the **Related Alerts** tab.

On the right, you can see the following:
- If the same alert was triggered on the object in the past 30 days. This helps you understand if this is a recurring problem or something new.
- If the same alert was triggered on other peers in the same environment, in the past 30 days. This helps you do a quick peer analysis to understand if others are impacted with the same problem.
- All the alerts triggered in the current topology. This helps you investigate if there are other alerts upstream or downstream in the environment which are impacting the health of the object.

**Potential Evidence Tab**

See the **Potential Evidence** tab for potential evidences around the problem, and to arrive at the root cause. This tab displays events, property changes, and anomalous metrics potentially relevant to the alert. The time range and the scope are fixed. To modify the scope or the time range and investigate further, click **Launch Workbench**. This runs the troubleshooting workbench.

The time range that is displayed in the potential evidence tab is two hours and thirty minutes before the alert was triggered. VMware Aria OperationsVMware Cloud Foundation Operations looks for potential evidences in this time range.

**Symptoms Display**

The symptoms display includes all the symptoms triggered for the current object. Use the symptom list to identify problems with an object so that you can resolve alerts generated for the object.

**How the Symptoms Work**

The list is the active triggered symptoms for an object, either as part of a generated alert or as a triggered symptom that is not included in an alert. This complete symptom list is useful for identifying problems that occur on an object but are not currently included in your alert definitions.

Click a symptom in the list to display the symptom details. An arrow in each column heading enables you to order the list in ascending or descending order. You can select multiple rows in the list using Shift+click, Control+click.

**Where You Find the Symptoms Display**

- From the left menu, click **Global Inventory**, then select a group, application, or inventory object. Click the **object** to display the object's **Summary** tab. Click the **Symptoms** subtab in the **Alerts** tab.
- In the menu, select **Search** and locate the object of interest. Click the **object** to display the object's **Summary** tab. Click the **Symptoms** subtab in the **Alerts** tab.

**Table 365: Symptoms Data Grid**

| Option | Description |
|---|---|
| Criticality | Criticality is the level of importance of a symptom in your environment.<br><br>The level is based on the same level assigned when the symptom was created. The possible values include:<br>• Critical<br>• Immediate<br>• Warning<br>• Information |
| Symptom | Name of the triggered symptom. |
| Status | Current state of the symptom.<br><br>Possible values are Active or Inactive. |
| Created On | Date and time when the alert was generated. |
| Canceled On | Date and time when the symptom was canceled. |
| Information | Information about the triggering condition for the symptom, including the trend and current value.<br><br>The sparkline displays a range of data that includes six hours before the symptom update time and one hour after the update time. |

**Table 366: Actions Menu**

| Option | Description |
|---|---|
| Actions menu | Select the **Go To Symptom Definition** option to open the **Symptom Definitions** page in the **Configure › Alerts** path. |

**Table 367: Include**

| Option | Description |
|---|---|
| Self | Shows or hides symptoms for the current object. |
| Peer | Shows or hides symptoms for objects like the impacted object. |
| Parents <options> | Shows or hides symptoms for the parent, grandparent, and so on, objects of the current object. |
| Children <options> | Shows or hides the symptoms for the descendants of the impacted object. |

**Table 368: Filters**

| | |
|---|---|
| Filtering options | The advanced search and filter lets you search for a symptom by:<br>• Symptom<br>• Status<br>• Criticality<br>• Triggered On<br>• Created On<br>• Canceled On |

**Table 369: Group By Options**

| Option | Description |
|---|---|
| None | Symptoms are not sorted into specific groupings. |
| Time | Group Symptoms by time triggered. |
| Criticality | Group Symptoms by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. |
| Definition | Group Symptoms by definition, that is, group like Symptoms together. The default. |
| Object Type | Group Symptoms by the type of object that triggered the alert. For example, group Symptoms on hosts together. |
| Scope | Group Symptoms by scope. You can search for Symptoms within the selected scope. |

### Events Display

An event is any change to an object defined by a change in the metrics for that object. You can compare changes to an object with symptoms and other data to identify a possible cause for a generated alert.

#### How the Events Display Works

The Events tab is a list of all the events generated for the selected object, group, or application based on messages send by either external sources or by internal self troubleshooting system.

The following vCenter activities are some of the activities that generate VMware Aria OperationsVMware Cloud Foundation Operations events:

- Powering a virtual machine on or off
- Creating a virtual machine
- Installing VMware Tools on the guest OS of a virtual machine
- Adding a newly configured ESX/ESXi system to a vCenter system

Depending on alert definitions, these events might generate alerts.

You might monitor the same virtual machines with other applications that provide information to VMware Aria OperationsVMware Cloud Foundation Operations, with the adapters for those applications configured to provide change events. In this instance, the **Events** tab includes certain change events that occur on the monitored objects. These change events might provide further insight into the cause of problems that you are investigating.

**Where You Find the Events Display**

- From the left menu, click **Global Inventory**, then select a group, custom data center, application, or inventory object. Click the **object** to display the object's **Summary** tab. Click the **Events** subtab in the **Alerts** tab.
- In the menu, select **Search** and locate the object of interest. Click the **object** to display the object's **Summary** tab. Click the **Events** subtab in the **Alerts** tab.

**Table 370: Include**

| Option | Description |
|---|---|
| Self | Shows or hides events for the current object. |
| Peer | Shows or hides events for objects like the impacted object. |
| Parents <options> | Shows or hides events for the parent, grandparent, and so on, objects of the current object. |
| Children <options> | Shows or hides the events for the descendants of the impacted object. |

**Table 371: Events Data Grid**

| Option | Description |
|---|---|
| Criticality | Criticality is the level of importance of the event in your environment. The event criticality appears in a tooltip when you hover the mouse over the criticality icon.<br><br>The level is based on the level assigned when the event definition was created, or on the highest symptom criticality, if the assigned level was **Symptom Based**.<br><br>The possible values include:<br>• Critical<br>• Immediate<br>• Warning<br>• Information |
| Object Name | Name of the event definition that generated the event. Click the event name to view the event details tabs where you can begin troubleshooting the event. |
| Event Type | Describes the type of event that triggered on the selected object, and helps you categorize the event so that you can assign certain types of events to specific system administrators. For example, Application, Virtualization/ Hypervisor, Hardware, Storage, Network, Administrative, and Findings. |
| Information | Describes additional information about the type of event that triggered on the selected object, and helps you categorize the events to a more detailed level. |
| Created On | Date and time when the event was generated. |
| Cancelled On | Date and time when the event was cancelled. |

**Table 372: Filters**

| | |
|---|---|
| Filtering options | The advanced search and filter lets you search for a event by:<br>• Symptom<br>• Status<br>• Criticality<br>• Triggered On<br>• Created On<br>• Canceled On |

**Table 373: Group By Options**

| Option | Description |
|---|---|
| None | Events are not sorted into specific groupings. |
| Time | Group events by time triggered. |
| Criticality | Group events by criticality. Values are, from the least critical: Info/Warning/Immediate/Critical. |
| Definition | Group events by definition, that is, group like events together. |
| Object Type | Group events by the type of object that triggered the alert. For example, group events on hosts together. |
| Scope | Group events by scope. You can search for events within the selected scope. |

### Events Timeline Tab

The generated alerts, triggered symptoms, and change events for the current object over time appear on the **Timeline** tab. You use the timeline to identify common trends over time that are contributing to the status of objects in your environment.

#### How the Events Timeline Works

The timeline view includes alerts, symptoms, and events for the selected object for the last 6 hours. To view the data for a particular time, click the timeline in one of the three tiers. Then move your mouse to the left to see data from the past or to the right to move back to the present.

The view is limited to approximately 50 alerts, symptoms, and events. If your timeline includes more than this number, you can use the toolbar options to remove data from the timeline until it contains data that you find useful for your investigation.

#### Where You Find the Events Timeline

- From the left menu, click **Environment**, then select a group, custom data center, application, or inventory object. Click the **object** to display the object's **Summary** tab. Click the **Events > Timeline** tabs.
- In the menu, select **Search** and locate the object of interest. Click the **object** to display the object's **Summary** tab. Click the **Events > Timeline** tabs.

**Table 374: View From Menu**

| Option | Description |
|---|---|
| Self | Shows or hides events for the current object. |
| Peer | Shows or hides events for objects like the impacted object. |
| Parents <options> | Shows or hides events for the parent, grandparent, and so on, objects of the current object. |
| Children <options> | Shows or hides the events for the descendants of the impacted object. |

**Table 375: Alert Filters**

| Option | Description |
|---|---|
| Criticality <options> | Limits the alerts to those matching the selected criticality level. If no criticality is selected, all alerts are displayed. |
| Status <options> | Limits the alerts in the chart to the canceled or active alerts. If no status is selected, all alerts are displayed.<br><br>This option applies only to alerts, not to fault and change events. Change events and active faults are always displayed in the chart. |
| Alert Type <options> | Select one or more alert types. The types are assigned when the alert is defined. If no type is selected, all alerts are displayed. |

**Table 376: Event Filters**

| Option | Description |
|---|---|
| Dynamic Threshold Violation | VMware Aria OperationsVMware Cloud Foundation Operations calculates dynamic thresholds for each metric that is collected for an object based on policies set. |
| Hard Threshold Violation | Events that represent a hard threshold violation, based on policies set. The system analyses the number of metrics that are violating their hard thresholds to determine trends. |
| Data Availability | Events reflecting datastore performance. Data availability is the capacity to provide data on demand to users and applications. |
| System Degradation | Events that reflect negative impacts on system performance. |
| Environment | Events indicating a change in the environment. |
| Change | Shows or hides the change events. Change events are changes to the object that might or might not result in an alert. |
| Notification | Routine notification events. |
| Fault | Events indicating any observed behavior that differs from the expected one. |

**Table 377: Date Controls, Data Values, Events Chart**

| Option | Description |
|---|---|
| Date Controls | Limits the data in the chart to the selected time frame. |
| Data Values | When you click a data point, the event is highlighted in the event data grid. |
| Events chart | Shows the events and alerts over time by criticality, and other data options you select in the toolbar. |

### Topology Tab

The Topology tab allows you to view the relationship of the objects in a topological view.

**Table 378: Topology View Toolbar Options**

| Option | Description |
|---|---|
| Back to initial object | Returns the map to original object if you double-clicked on an icon to examine another object. |
| Vertical/Horizontal | Displays a vertical or horizontal view of the graph or tree view. |
| Hide Text/Show Text | Hides or displays the object names. |
| Standard View/Fit View | **Standard View**: Fixes the view to a specific zoom level.<br>**Fit View**: Adjusts the graph or tree view to fit the screen. |
| Settings | • **Parent Depth**: Displays up to five levels of parent objects for the selected object depending on your selection.<br>• **Child Depth**: Displays up to five levels of child objects for the selected objects, depending on your selection.<br>• **Page Size**: Sets the number of objects per page. |

**Table 379: Specific Functions from the Topology Grid View**

| | |
|---|---|
| Hide Node | Hover your cursor over an object and click the **Hide Node** icon to hide the specific node. |
| Show Peers/Hide Peers | Hover your cursor over the selected object in the topology pane and select **ShowPeers/Hide Peers** to view or hide other objects of the same type that exist under the parent object. |
| Pagination/Next Page/Previous Page | At the bottom of the topological view you can view the page number and also move to the next/previous page. |
| Filter | At the bottom of the topological view you can use the filter option to search for objects. |
| Arrows | Use the arrows on each object to view relationships of each object. |

### Creating and Using Object Details

The Views and List tabs provide you with specific data about the object. You use this information to evaluate problems in more detail and to view descendant objects.

### Details Views Tab

The **Views** tab is divided into two panels. The bottom panel is updated depending on the view you select in the top panel.

In the top panel you can create, edit, delete, clone, export, and import views. The views list depends on the object you select from the environment. Each view is associated with an object. For example, the predefined VM inventory - Memory list view is available when you select a host.

You can limit the views list by adding a filter from the right side of the panel. Each of the provided filter groups limits the list by the word you type.

The Filters option helps you to list views according to the following criteria:
- Name
- Type
- Description
- Subject
- Owner

For example, if you select **Description** and type `my view`, the listed views are all views that are applicable for the selected object and contain *`my view`* in the description.

**Table 380: Views List Table Columns**

| Column | Description |
|---|---|
| Name | Name of the view. |
| Type | Type of the view. A view type is the way the collected information for the object is presented. |
| Description | Description of the view as it is defined when the view is created. |
| Subject | Object type with which a view is associated. |
| Owner | Owner of the view is the user, who created it or edited it for the last time. |

In the bottom panel of the **Views** tab, you can see the data of the object, calculated by a selected view from the top panel. Say, for example, the selected object is a host and you select Virtual Machine Configuration Summary List View. The result is a list of all the virtual machines on that host, and their data calculated by the view.

For Trend views, you can select a parent object and see the data of the associated child objects and metrics in the bottom panel of the **Views** tab.

For Distribution views, you can click on a section of the pie chart or on one of the bars in the bar chart to view the list of objects filtered by the selected segment, in the bottom panel of the **Views** tab.

#### Where You Find the Details View Tab

- From the left menu, click **Inventory**, then select an inventory object from the **Inventory Detailed View** panel. Click the **Details** tab, then select the **Views** tab.

## Details List Tab

VMware Aria OperationsVMware Cloud Foundation Operations discovers the descendants of the selected object and lists them.

### How the List Tab Works

Descendant objects of the object you selected appear in the datagrid of the **List** tab. To find a particular object, you can sort a column in the grid or search for a word in the filter, or use the filter options.

### Where You Find the List Tab

From the left menu, click **Inventory** and then click a vCenter Server instance from the datagrid. From the **Inventory** panel, select an object and then click the **Details › List** tab from the right side.

Use the toolbar options to manage objects.

- Filter options limit the list to objects matching the filter. Filter options include All, ID, Internal ID, Name, Description, Adapter Type, Object Type, and Identifiers.
- Select the object to manage from the list.

**Table 381: Toolbar Options**

| Option | Description |
|---|---|
| Actions | Perform an action on the selected object. Available actions depend on the object type. For example, Power on VM applies to the selected virtual machine. |
| Open in external application | If an adapter includes the ability to link to another application for information about the object, click the button to access a link to the application. For example, Open Virtual Machine in a vSphere Client or Search for VM logs in VMware Cloud Foundation Operations for logs. |
| Edit Object | Edit the selected object. For example, add or change the maintenance schedule for a virtual machine. If multiple objects of the same type are selected, common identifiers for the object type are editable. For example, change the VM entity name of multiple datastores with a single edit. See Manage Objects Workspace. |
| Add Object | VMware Aria OperationsVMware Cloud Foundation Operations discovers objects for most adapters. For adapters that do not support autodiscovery for all objects, the objects are manually added. See Manage Objects Workspace. |
| Discover Objects | Perform an IP scan to discover objects associated with a particular adapter. See Discover Objects Workspace. |
| Delete object | Remove the object from the list. |
| Start maintenance | Take the object offline for maintenance. See Manage Maintenance Schedules for Your Object Workspace. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| End maintenance | Terminate the maintenance period and put the selected object back online. |
| Clear Selections | Clear all object selections. |
| Select All | Select all objects displayed. |
| Add/Edit Custom Property | Define or add custom properties. |
| Go to Details | Display the **Summary** tab of the selected object. |
| Page Size | The number of objects to list per page. |

**Table 382: Datagrid Options**

| Option | Description |
|---|---|
| Name | Displays the name of the object. |
| Adapter Type | Displays the name of the adapter. |
| Object Type | Displays the type of object. |
| Policy | Displays the policy associated with the object. |
| Collection State | Displays the collection state of an adapter instance for each object. |
| Collection Status | Displays the collection status of an adapter instance of each object. |

### User Scenario: Investigate the Root Cause of a Problem by Using the Troubleshooting Tab Options

One of your customers reports poor performance for a virtual machine, including slowness and fails. This scenario provides one way that you can use VMware Aria OperationsVMware Cloud Foundation Operations to investigate the problem based on information available in the **Troubleshooting** tabs.

As a virtual infrastructure administrator, you respond to a help ticket in which one of your customers reports problems with a virtual machine, sales-10-dk. The reported conditions are poor application performance, including slow load times and slow boot, some applications are taking longer and longer to load, and files are taking longer to save. Today applications started to fail and an update failed to install.
When you look at the **Alerts** tab for the virtual machine, you see an alert for chronic high memory workload leading to memory stress. The triggered symptoms indicate memory stress and the recommendation is to add more memory.
Based on experience, you are not convinced that this alert indicates the root cause, so you review the **Capacity** tab. The **Capacity** tab indicates memory and disk space problems, and Time Remaining, which has 0 days remaining for memory and disk space.
From this initial review, you know that problems exist in addition to the memory alert, so you use the **Events** tabs to do a more thorough investigation.

### Review the Triggered Symptoms When You Troubleshoot a Virtual Machine Problem

As a virtual infrastructure administrator, you respond to customer complaints and alerts, and identify problems that occur on the objects in your environment. You use the information on the **Symptoms** tab to help determine whether the triggered symptoms indicate conditions that contribute to the reported or identified problem.

You must research a problem of poor performance on one of your virtual machines, as reported by one of your customers. When you view the **Alerts** tab for the virtual machine, the only alert that appears is named `Virtual Machine is Violating Risk Profile 1 in vSphere Hardening Guide`.
When you reviewed the **Capacity** tab for the virtual machine, you identified that problems were occurring with memory and disk space. Now, you focus your attention to the triggered symptoms on the virtual machine.

The following method of using the **Symptoms** tab to evaluate problems is provided as an example for using VMware Aria OperationsVMware Cloud Foundation Operations , and is not definitive. Your troubleshooting skills and your knowledge of the particular aspects of your environment determine which methods work for you.

1. In the menu, click **Dashboards**, then click **Troubleshoot a VM** in the left pane.
2. Search for a virtual machine to troubleshoot.
   In this example, the virtual machine name is named `sales-10-dk`.
3. With the virtual machine selected, click the **Alerts** tab, and click the **Symptoms** tab.
4. Review and evaluate the triggered symptoms.

| Option | Evaluation Process |
|---|---|
| Symptom | Are any of the triggered symptoms related to the critical states you see for memory or disk space? |
| Status | Are the symptoms active or inactive? Even inactive symptoms can provide information about the past state of the object. To add any inactive symptoms, click **Status: Active** on the toolbar to remove the filter. |
| Created On | When did the symptoms trigger? How does the time of the triggered symptom compare with the other symptoms? |
| Information | Can you identify a correlation between the triggered symptoms and the state of the Time Remaining and Capacity Remaining badges? |

From your review, you determine that some of the triggered symptoms are associated with compliance alerts for the virtual machine as defined in the *`vSphere Hardening Guide`*. The violated symptoms triggered for the alert named *`vSphere Hardening Guide`*, which is one of several compliance risk profiles provided with VMware Aria OperationsVMware Cloud Foundation Operations .
The following symptoms triggered in the compliance alert named `Virtual Machine is Violating Risk Profile 1 in vSphere Hardening Guide`:

- `Independent nonpersistent disks are being used`
- `Autologon feature is enabled`
- `Copy/paste operations are enabled`
- `Users and processes without privileges can remove, connect and modify devices`
- `Guests can receive host information`

Other symptoms also triggered, which are related to memory and time remaining.

- `Guest file system overall disk space usage reaching critical limit`
- `Virtual machine disk space time remaining is low`
- `Virtual machine CPU time remaining is low`
- `Guest partition disk space usage`
- `Virtual machine memory time remaining is low`

Review the symptoms for the object on a timeline. See Compare Symptoms on a Timeline When You Troubleshoot a Virtual Machine Problem.

You can find the *`vSphere Hardening Guides`* at http://www.vmware.com/security/hardening-guides.html.

### Compare Symptoms on a Timeline When You Troubleshoot a Virtual Machine Problem

Looking at the triggered symptoms for an object over time helps you to compare triggered symptoms, alerts, and events when you are troubleshooting problems with objects in your environment. The **Timeline** tab in VMware Aria OperationsVMware Cloud Foundation Operations provides a visual chart on which to see triggered symptoms that you can use to investigate problems in your environment.

Review the triggered object symptoms. See Review the Triggered Symptoms When You Troubleshoot a Virtual Machine Problem.

After you identify the following symptoms as possible indicators of the root cause of the reported performance problems on the sales-10-dk virtual machine, you compare them to each other over time. Look for unusual or common patterns.

- Guest file system overall disk space use reaching critical limit.
- Virtual machine disk space time remaining low.
- Virtual machine CPU time remaining low.
- Guest partition disk space use.
- Virtual machine memory time remaining is low.

The following method of evaluating problems using the **Timeline** tab is provided as an example for using VMware Aria OperationsVMware Cloud Foundation Operations and only one method. Your troubleshooting skills and your knowledge of the specifics of your environment determine which methods work for you.

1. Enter the name of the virtual machine in the **Search** text box on the main title bar.
   In this example, the virtual machine name is `sales-10-dk`.
2. Click the **Events** tab and click the **Timeline** tab.
3. On the Timeline toolbar, click **Date Controls** and select a time that is on or before the reference symptoms were triggered.

   The default time range is the last 6 hours. For a broader view of the virtual machine over time, configure a range that includes triggered symptoms and generated alerts.
4. To view the point at which the symptoms were triggered and to identify which line represents which symptom, drag the timeline week, day, or hour section left and right across the page.
5. Click **Event Filters** and select all the event types.

   Consider whether events correspond to triggered symptoms or generated alerts.
6. In the Related Hierarchies list in the upper left pane, click **vSphere Hosts and Clusters**.

   The available ancestors and descendant objects depend on the selected hierarchy.
7. To see if the host is experiencing a contributing problems, click **View From** and select **Host System** under Parent.

   Consider whether the host has symptoms, alerts, or events that provide you with more information about memory or disk space problems.

Comparing virtual machine symptoms to host symptoms, and looking at the symptoms over time indicates the following trends:

- The host resource use, host disk use, and host CPU use symptoms are triggered for about 10 minutes approximately every 4 hours.
- The virtual machine guest-file system out-of-space symptom is triggered and canceled over time. Sometimes the symptom is active for an hour and canceled. Sometimes it is active for two hours. But no more than 30 minutes occur between cancellation and the next triggering of the symptom.

Look at events in the context of the badges and alerts. See Identify Influential Events When You Troubleshoot a Virtual Machine Problem.

### Identify Influential Events When You Troubleshoot a Virtual Machine Problem

Events are changes to objects in your environment that are based on changes to metrics, properties, or information about the object. Examining the events for the problematic virtual machine in the context of alerts can provide visual clues to the root cause of a problem.

Examine triggered symptoms, alerts, and events over time. See Compare Symptoms on a Timeline When You Troubleshoot a Virtual Machine Problem.

As a virtual infrastructure administrator investigating a reported performance problem with a virtual machine, you compared symptoms on the timeline. You identified odd behavior related to a guest file system that you want to examine in the context of other metrics. This investigation can determine whether you find the root cause of the problem.
The following method of evaluating problems using the **Events** tab is provided as an example for using VMware Aria OperationsVMware Cloud Foundation Operations and is not definitive. Your troubleshooting skills and your knowledge of the particulars of your environment determine which methods work for you.

1. Enter the name of the virtual machine in the **Search** text box, on the main title bar.
   In this example, the virtual machine name is sales-10-dk.
2. Click the **Events** tab and select the **Events** button.
3. On the Events toolbar, click **Date Controls** and select a time that is on or before the symptoms were triggered.
4. Click **Event Filters** and select all the event types.

   Consider whether any changes correspond to other events.

5. Click **View From > Parent › Select All** and click through the alerts in the timeline to review events.

   Consider whether any of the events, which are listed in the data grid below the chart, correspond to problems with the host that might contribute to the reported problem.

6. Click **View From > Child › Select All** and click through the alerts to review the events.
   Consider whether any of the events show problems with the datastore.

Your evaluation shows no particular correlation between the workload and the time at which the guest file system out-of-space symptom was triggered each time.

### Synthetic Monitoring Tab

The Synthetic Monitoring tab displays the API call monitoring details for business applications that have synthetic monitoring configured and enabled.

The Synthetic Monitoring tab has four panes. The cards in the top pane display the following information:
- Summary Card - displays a summary of the Synthetic Monitoring configuration.
- Business Application KPI - Displays the total availability of the all application API calls averaged month to date.

The middle pane displays an overview of the performance status of all the APIs configured along with the average response time. You can select a custom period, or choose preset time periods such as 1 hour, 6 hours, 24 hours, or 7 days.

In the bottom most pane, you can see the status of all the configured API calls with the option of expanding each call. You can view the following information:
- Total Response Time
- DNS Lookup
- TCP Connection

- TLS Handshake
- Server Processing
- Content Transfer

When you define alerts using the metrics or properties which are pushed by Synthetic Monitoring, you can observe the symptom description from the Synthetic Monitoring tab when expanding the API Request row. The Performance on API Calls chart is colored using the criticality of the triggered symptom which is defined on the Response Time metric.

## Running Actions from VMware Aria OperationsVMware Cloud Foundation Operations

The actions available in VMware Aria OperationsVMware Cloud Foundation Operations allow you to modify the state or configuration of selected objects in vCenter from VMware Aria OperationsVMware Cloud Foundation Operations. For example, you might need to modify the configuration of an object to address a problematic resource issue or to redistribute resources to optimize your virtual infrastructure.

The most common use of the actions is to solve problems. You can run them as part of your troubleshooting procedures or add them as a resolution recommendation for alerts.

When you grant a user access to actions in VMware Aria OperationsVMware Cloud Foundation Operations, that user can take the granted action on any object that VMware Aria OperationsVMware Cloud Foundation Operations manages.

When you are troubleshooting problems, you can run the actions from the center pane Actions menu. Alternatively, you can run them from the toolbar on list views that contain the supported objects.

When an alert is triggered, and you determine that the suggested action is the most likely way to resolve the problem, you can run the action on one or more objects.

**Run Actions from Toolbars in VMware Aria OperationsVMware Cloud Foundation Operations**

When you run actions in VMware Aria OperationsVMware Cloud Foundation Operations, you change the state of vCenter objects. You run one or more actions when you encounter objects where the configuration or state of the object is affecting your environment. These actions allow you to reclaim wasted space, adjust memory, or conserve resources.

- Verify that the vCenter Adapter is configured to run actions for each vCenter instance. See Configure a vCenter Serve Cloud Account in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* .
- Verify that the vCenter Adapter is configured to run actions for each vCenter instance. See the *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* .
- Ensure that you understand how to use the power-off-allowed option if you are running Set CPU Count, Set Memory, and Set CPU Count and Memory actions. See Working with Actions That Use Power Off Allowed section in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* .
- Ensure that you understand how to use the power-off-allowed option if you are running Set CPU Count, Set Memory, and Set CPU Count and Memory actions. See the section Working With Actions That Use Power Off Allowed in the VMware Aria OperationsVMware Cloud Foundation Operations Information Center.

This procedure for running actions is based on the VMware Aria OperationsVMware Cloud Foundation Operations**Actions** menus and is commonly used when you are troubleshooting problems. The available actions depend on the type of objects with which you are working. You can also run actions as alert recommendations.

1. Select the object in the Environment page inventory trees or select one or more objects it in a list view.
2. Click **Actions** on the main toolbar or in an embedded view.
3. Select one of the actions.
   If you are working with a virtual machine, only the virtual machine is included in the dialog box. If you are working with clusters, hosts, or datastores, the dialog box that appears includes all objects.
4. To run the action on the object, select the check box and click **OK**.
   The action runs and a dialog box appears that displays the task ID.

5. To view the status of the job and verify that the job finished, click **Recent Tasks** or click **OK** to close the dialog box. The Recent Tasks list appears, which includes the task you just started.

To verify that the job completed, click **Environment** in the menu and click **History >Recent Tasks**. Find the task name or task ID in the list and verify that the status is finished. See Monitor Recent Task Status.

### Rebalance Container Action

When the workload in your environment becomes imbalanced, you can move the workload across your objects to rebalance the overall workload. The container for the rebalance action can be a data center or a custom data center, and the objects that are moved are the virtual machines in the suggested list provided by the action.

#### DRS Must be Activated on Clusters

Your vCenter instance must have a cluster that passes a DRS-activated check for the Rebalance Container action to appear in the Actions drop-down menu.

To get the Rebalance Container action from a custom data center or data center, and the related alerts, you must have the following:

- A vCenter Adapter configured with the actions activated for each vCenter instance
- A vCenter instance with at least one cluster that is DRS-activated.

If your cluster does not have DRS fully automated, the Rebalance Container action notifies you that one or more clusters under the selected container do not have DRS set to fully automated.

To ensure that the Rebalance Container action is available in your environment, you must add DRS. Then, wait one collection cycle for the Rebalance Container action to appear.

#### You Must Have Access to All Objects in the Container

If you have access to all objects in a cluster, data center, or custom data center, you can run the Rebalance Container action to move virtual machines to other clusters. When you do not have access to all of the objects in the container, the Rebalance Container action is not available.

#### How the Rebalance Container Action Works

If two data centers are experiencing extreme differences in workload - one high and one low - use the Rebalance Container action to balance the workload across those objects. For example, if the CPU demand on a host in one data center exceeds its available CPU capacity, critical pressure occurs on the host. To identify the cause of stress, monitor the CPU demand. Some virtual machines on each host might be experiencing high CPU demand, whereas others might be experiencing a low demand.

The Rebalance Container action moves all affected objects in the suggested list provided by the action to balance the workload. If you do not want to act on the entire set of objects to resolve the problem with workload, you can use the Move VM action to move an individual object.

> **IMPORTANT**
> Do not attempt to move virtual machines that are members of a vApp, because the vApp can become nonfunctional. Instead, add affinity rules for these virtual machines to keep them together so that the Move VM and Rebalance Container actions will ignore them.

When workloads become imbalanced, the following alerts can trigger on data centers and custom data centers. These alerts are deactivated by default in the policies.

- Custom data center has unbalanced workload
- Data center has unbalanced workload

When the workloads on hosts in a data center or custom data center differ significantly, click **Operations › Alerts** and verify whether the alert triggered. For example, to verify whether the alert triggered on a custom data center, check the alert named `Custom data center has unbalanced workload`. You can click the alert to view the causes of the alert and identify the source of the imbalance problem on the **Summary** tab.

To display the recommendations about the objects to move so that you can rebalance the workload, click the **Rebalance Container** action on the **Summary** tab. The recommendations indicate that you move one or more virtual machines to another host. When you click **OK**, a pop-up message provides a link to track the status of the action in **Recent Tasks**.

The action moves the virtual machines identified in the recommendation to the host machine that has a low workload or stress. You can view the status of the action in the list of recent tasks in **Administration › Control Panel**, and then click the **Recent Tasks** tab. You can also use the vSphere Web Client to view the status of the action and the performance for the host.

After the action runs and VMware Aria OperationsVMware Cloud Foundation Operations performs several collection cycles, view the workload on the data center to confirm that the workload was rebalanced and that the alert is gone.

## Where You Run the Action

You can run the Rebalance Container action from the Actions menu for a data center or custom data center, or you can provide it as a suggested action on an alert.

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- From the left menu click **Inventory**, select an object, click the **Details** tab, click **Views**, and select a view of type List.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and select an object in the **List** tab.
- From the left menu click **Operations › Configurations**, then click the **Objects** tab, and select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

## Action Recommendations

Review the following information about the hosts and virtual machines to ensure that you are submitting the action for the correct objects.

| Option | Description |
|---|---|
| Virtual Machine | Name of the virtual machine on the host that is experiencing an excessive workload. |
| Source Cluster | Name of the cluster on which the virtual machine is running. |
| Datastores | Datastore associated with the virtual machine. |
| Destination Cluster | Cluster where the virtual machine is to be moved. DRS selects the host automatically. |
| Reason | Describes the action to be taken and the reason why the move is suggested. For example, the recommendation is to |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | move part of the workload on the cluster to another cluster to reduce the imbalance in CPU demand. |
| Parent vCenter | Identifies the vCenter Server adapter associated with the affected cluster. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 383: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Delete Idle VM Action

The Delete Idle VM action in VMware Aria OperationsVMware Cloud Foundation Operations removes from your vCenter instances those selected virtual machines that are in an idle state. Use this action to reclaim redundant resources.

#### How the Action Works

The Delete Idle VM action removes from your vCenter instances those virtual machines that are powered on, but that are in an idle state.

#### Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu, click **Inventory**, select an object, click the **Details** tab, and click **Views**.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

#### Action Menu Items

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

| Menu Items | Description |
|---|---|
| Name | Name of the virtual machine as it appears in the environment inventory. |
| Host | Name of the host on which the virtual machine is running. |
| Parent vCenter | Parent vCenter instance where the virtual machine resides. |

After you click **Begin Action**, the next dialog box provides the task ID and a link to the task list.

**Table 384: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Set DRS Automation Action

You can monitor and configure the vSphereDistributed Resource Scheduler (DRS) automation rules from VMware Aria OperationsVMware Cloud Foundation Operations. DRS monitors and allocates the resources in your environment, and balances the computing capacity across your hosts and virtual machines.

#### How the Action Works

The Set DRS Automation action monitors and configures DRS automation rules. With the Set DRS Automation action, you can activate and deactivate DRS.

If VMware Aria Automation manages any of the virtual machines in your environment, the Set DRS Automation action is not available for that object.

#### Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and click **Views**.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

#### Action Menu Items

To ensure that you are submitting the correct action for the correct objects, review the following information about the clusters.

| Menu Items | Description |
|---|---|
| Name | Name of the cluster in the vCenter instance. |
| Automation Level | Level of DRS automation. When DRS is fully automated on the selected cluster, you can run the Set DRS Automation action. |
| Migration Threshold | Recommendations for the migration level of virtual machines. Migration thresholds are based on DRS priority levels, and are computed based on the workload imbalance metric for the cluster. |
| Parent vCenter | Parent vCenter instance where the cluster resides. |

After you click **Begin Action**, the next dialog box provides the task ID and a link to the task list.

**Table 385: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Execute Script Action

To troubleshoot particular processes, you can upload a script or run a command to receive specific information. You can view the standard output or standard error as applicable.

#### Where You Run the Action

For supported objects and object levels, in the main menu, select the **Inventory** tab and then select the relevant VM from the Inventory tree. This action is available from the **Actions** menu just below the top menu in VMware Aria OperationsVMware Cloud Foundation Operations.

#### Prerequisites

- VMware Tools must be installed and running on the VM. For details see KB 75122
- Service discovery is activated with the successful discovery of VMs.
- The VM must be powered on and connected.

#### Action Options

Enter the VM credentials to authenticate even when the VM guest OS authentication status is "Success". You can run a script by entering it directly or by uploading a script file by optionally providing arguments.

| Option | Description |
|---|---|
| Upload File | Use this option to browse and upload the script that you want to run. |
| File | Browse and upload the script file. |
| Args | List the arguments in the script. |
| Command | Select the option and enter a command in the text box. |
| Timeout | Script execution timeout on VMs. Script execution continues even if the dialog box is closed. You can verify the status from **Recent Tasks**. |
| Execute | Runs the script or command. |
| stdout | Displays the standard output. |
| stderr | Displays errors, if any. |

### Get Top Processes Action

The Get Top Processes action is used for troubleshooting process issues and resource issues related to the applications of the virtual machine.

### How the Action Works

The Get Top Processes action, provides the status of top 10 processes for the selected virtual machine. You can troubleshoot issues related to the resources that are affecting the applications in the virtual machine.

By default, the details of top 10 processes are displayed for the selected virtual machine. You can change the number of processes and view the details for top N processes where N is between 1-100. You have the option to view the processes based on CPU and Memory.

The Get Top Processes action is run on both Windows virtual machine and Linux virtual machine. You can view the summary information for the commands only in a Linux virtual machine.

### Where You Run the Action

For supported objects and object levels, in the main menu, select the **Environment** tab and then select the relevant VM from the Inventory tree. This action is available from the **Actions** menu just below the top menu in VMware Aria OperationsVMware Cloud Foundation Operations.

### Prerequisites

- VMware Tools must be installed and running on the VM. For details see KB 75122
- Service discovery is activated with the successful discovery of VMs.
- The VM must be powered on and connected.

### Action Options

You must enter the VM credentials to authenticate when the VM is monitored in a credential-less mode or when the VM is monitored in a credential-based mode where the user is not authenticated. To ensure that you are taking the right action, review the following information.

| Option | Description |
|---|---|
| Number of Processes | Displays the number of processes for which the details are displayed. |
| Refresh | Displays new data about processes, when you change the value for the number of processes. |
| Command | Displays the name of the application |
| PID | Displays the process ID. |
| CPU | Displays the CPU usage in percentage for Linux VMs.<br><br>Displays the CPU usage in seconds for Windows VMs. The count starts when you start the operating system in the VM . |
| Mem (%) | Displays the Memory usage in KB. |
| User | Displays the user name. |
| Status | Displays the process status. It can be in one these states:<br>• For Linux - I, R, S<br>• For Windows - Unknown, Running, and Sleeping |
| Run | Displays data about the specified numbers of processes. |

## Move Virtual Machine Action

You can use the Move VM action to move virtual machines from one host and datastore to another host and datastore to balance the workload in your environment.

### How the Action Works

When you initiate this action, the **Move VM** wizard opens and scopes the possible destinations. You select the destination host and datastore from the list of available destinations.

To see all destinations, you must have view access to the following object types:

- Scope object, which includes a vCenter, data center, custom data center, or cluster.
- Host in the scope object.
- Datastore in the host.

The destinations include combinations of objects for the move, such as a specific host and datastore, or a different host with the same datastore. You select one of the available combinations. If your environment includes many destination objects, such as many hosts or datastores, enter text in the filter text box to search for specific destination objects.

The **All Filters** option helps you to move the VM according to the following action option:

- Destination Host
- Destination Datastore
- Will it Fit
- VM PowerOff Required
- Affinity Rules

VMware Aria OperationsVMware Cloud Foundation Operations uses vSphere DRS rules that you define in vCenter to help determine good placement decisions for your virtual machines in the move action. The Affinity Rules column indicates whether those rules are violated by the Move VM action.

> **IMPORTANT**
> Do not attempt to move virtual machines that are members of a vApp, because the vApp can become nonfunctional. Instead, add affinity rules for these virtual machines to keep them together so that the Move VM and Rebalance Container actions will ignore them.

To initiate the action, you click the **Begin Action** button.

When you finish the wizard, VMware Aria OperationsVMware Cloud Foundation Operations displays a dialog box to indicate that the action has started. To track the status of the action, click the link in the dialog box and view the state of the action in **Recent Tasks**.

### Moving Virtual Machines is Not Allowed Across Data Centers

When you attempt to use the **Move VM** action to move a virtual machine across data centers, VMware Aria OperationsVMware Cloud Foundation Operations must be able to identify the matching network and storage objects for the destination data center. Network objects include VMware virtual switches and distributed virtual switches. Storage objects include datastores and datastore clusters.

Moving a virtual machine across data centers requires VMware Aria OperationsVMware Cloud Foundation Operations to move the virtual machine files and change the virtual machine network configuration. VMware Aria OperationsVMware Cloud Foundation Operations does not currently move the virtual machine files across datastores, nor does it change the virtual machine network configuration. As a result, VMware Aria OperationsVMware Cloud Foundation Operations does not allow you to move virtual machines across data centers.

When you use the **Move VM** action, be aware of the following behavior:

- If you select a single virtual machine, VMware Aria OperationsVMware Cloud Foundation Operations displays the data center where the virtual machine resides.
- If you select multiple virtual machines, but those virtual machines do not share a common data center, the **Move VM** action does not display the data centers, and the **Move VM** action does not appear in the actions menu.

**Where You Run the Action**

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu, click **Inventory**, select an object, click the **Details** tab, and then click the **Views** tab.
- From the left menu, click **Inventory**, select an object, click the **Details** tab, and select an object from the **List** tab.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

**Action Options**

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

| Option | Description |
|---|---|
| Priority | Indicates the priority of the proposed move destination. When the action is automated, the proposed destination with priority of 1 is automatically selected. |
| Destination Host | Name of the host to which the virtual machine will be moved. |
| Current CPU Workload | Amount of CPU in GHz available on the host. |
| Current Memory Workload | Amount of memory in GB available on the host. |
| Destination Datastore | Datastore to which the virtual machines storage will be moved. |
| Current Disk Space Workload | Amount of disk space available on the datastore. |
| Will it fit | Calculated estimation of whether the virtual machine fits on the selected destination. |
| VM Power Off Required | When set to `No`, the action does not power off the virtual machine before the move. When set to `Yes`, the action powers off the virtual machine before the move takes place, and powers on the virtual machine after the move is complete. If VMware Tools is installed, a guest OS shutdown is used to power off the virtual machine. |
| Affinity Rules | Indicates whether vSphere DRS rules exist, as defined in vCenter. For example, a rule might exist to keep virtual machines together, and another rule might exist to separate virtual machines.<br><br>This column indicates the following status.<br><br>• Empty. vSphere DRS rules are not defined.<br>• Green check mark. The move of virtual machines does not violate affinity rules. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Red circle with bar. The move of virtual machines does break affinity rules. If you choose to break the affinity rules, you must resolve any problems manually. |
| Affinity Rule Details | Identifies the virtual machine and the vSphere DRS rule name as defined in vCenter. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 386: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Power Off Virtual Machine Action

The Power Off VM action in VMware Aria OperationsVMware Cloud Foundation Operations stops one or more selected virtual machines that are in a powered on state. You power off a virtual machine when you are managing resources and reclaiming wasted space.

#### How the Action Works

The Power Off VM action turns off the virtual machine. If VMware Tools is installed and running, the guest operating system is shut down before the machine is powered off. If VMware Tools is not installed and running, the virtual machine is powered off regardless of the state of the guest operating system. In this case, use this action only when you are powering off virtual machines where stopping the guest operating system does not adversely affect the installed applications.

If the target virtual machine is already powered off, the recent task status reports success on the machine, even though the state of the virtual machine did not change.

#### Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click **Views**.
- From the left menu click **Operations › Configurations**, and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

#### Action Options

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| Power State | Indicates whether the virtual machine is powered on or powered off. |
| Idle VM | Indicates whether the virtual machine is considered to be in the idle state based on the configured idle virtual machine metric.<br><br>Possible values include:<br>• false. The virtual machine is active.<br>• true. The virtual machine is idle.<br>• unknown. VMware Aria OperationsVMware Cloud Foundation Operations does not have the data required to calculate the idle metric. |
| Idle VM Percentage | Calculated threshold of the idle virtual machine percentage based on the configured reclaimable wasted space policy. |
| CPU Usage Percentage | Calculated threshold of the virtual machine CPU percentage based on the metric named `cpu \| usage_average`. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 387: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Shut Down Guest Operating System for Virtual Machine Action

The Shut Down Guest OS for VM action shuts down the guest operating system and powers off the virtual machine. You shut down a virtual machine when you are managing resources and reclaiming wasted space.

#### How the Action Works

The Shut Down Guest OS for VM action checks that VMware Tools, which is required, is installed on the target virtual machines, then shuts down the guest operating system and powers off the virtual machine. If VMware Tools is not installed or installed but not running, the action does not run and the job is reported as failed in **Recent Tasks**.

If the target virtual machine is already powered off, the recent task status reports success on the machine, even though the state of the virtual machine did not change.

**Where You Run the Action**

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and click **Views**.
- From the left menu click **Operations › Configurations**, and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

**Action Options**

Review the following so you can be sure you are taking the right action.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| Power State | Indicates whether the virtual machine is powered on or powered off. |
| Idle VM | Indicates whether the virtual machine is considered to be in the idle state based on the configured idle virtual machine metric.<br><br>Possible values include:<br>• false. The virtual machine is active.<br>• true. The virtual machine is idle.<br>• unknown. VMware Aria OperationsVMware Cloud Foundation Operations does not have the data required to calculate the idle metric. |
| Idle VM Percentage | Calculated threshold of the idle virtual machine percentage based on the configured reclaimable wasted space policy. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 388: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Reboot Guest OS for Virtual Machine Action

The Reboot Guest OS for VM action reboots the guest operating system and the virtual machine. You reboot a virtual machine while managing resources or when you have new updates or configuration changes to your virtual machine.

#### How the Action Works

The Reboot Guest OS for VM action checks that VMware Tools, which is required, is installed on the target virtual machines, then reboots the guest operating system and the virtual machine. If VMware Tools is not installed or installed but not running, the action does not run, and the job is reported as failed in **Recent Tasks**.

#### Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu, click **Inventory** and select an object in the list.
- From the left menu, click **Inventory**, select an object, click the **Details** tab, and click **Views**.
- From the left menu, click **Inventory**. Select an object, click the **Environment** tab, and select an object in the list view.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

#### Action Options

Review the following so you can be sure you are taking the right action.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| Host | Name of the host on which the virtual machine is running. |
| Parent vCenter | The adapter instance as configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 389: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Power on Virtual Machine Action

To start one or more virtual machines that are in a powered off state, use the Power On VM action. You power on a virtual machine so that you can shift resources. For example, power on a machine so that you can use it, run applications, or verify that actions that were run on already powered down machines contribute to improved performance.

#### How the Action Works

The Power On VM action powers on virtual machines that are powered off. The action does not affect virtual machines that are currently powered on.

If the target virtual machine is already powered on, the task status reports success for the machine even though the state of the virtual machine did not change.

#### Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and click **Views**.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click the **List** tab.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

#### Action Options

To ensure that you are taking the right action, review the following information .

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| Power State | Indicates whether the virtual machine is powered on or powered off. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 390: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Delete Powered Off Virtual Machine Action

The Delete Powered Off VM action in VMware Aria OperationsVMware Cloud Foundation Operations removes selected virtual machines that are in a powered off state from your vCenter instances. Use this action to reclaim redundant resources.

#### How the Action Works

The Delete Powered Off VM action removes virtual machines from the vCenter instances. If the virtual machine is powered on, the action does not delete the virtual machine.

#### Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click **Views**.
- From the left menu click **Operations › Configurations**, and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

#### Action Options

To ensure that you are submitting the action for the right objects, review the following information.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. |
| Name | Name of the virtual machine as it appears in the environment inventory. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Power State | Indicates whether the virtual machine is powered on or powered off. |
| Disk Space | Amount of disk space currently consumed by the virtual machine. |
| Snapshot Space | Amount of disk space currently consumed by the virtual machine snapshots. |
| Memory (MB) | Amount of memory allocated to the virtual machine. |
| CPU Count | Number of CPUs currently configured for the virtual machine. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 391: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Set Memory for Virtual Machine Action

The Set Memory for VM action in VMware Aria OperationsVMware Cloud Foundation Operations is used to add or remove memory on virtual machines. You increase the memory to address performance problems or decrease the memory to reclaim resources.

#### How the Action Works

The Set Memory for VM action perform several tasks. The action determines the power state of the target virtual machines, takes a snapshot when you request it and powers off the machine if necessary and you request it. As well, the action changes the memory to the new value, and returns the virtual machines their original power states.

An alternative form of the Set Memory for Virtual Machine action is available for automation. This action can run when the virtual machine is powered on or off.

Use this version of the action if the automated action has permission to power off the virtual machine, and hot add of memory is not activated on the virtual machine. With hot add activated, you can add memory, but you cannot remove it.

This version of the action would be required if a virtual machine is powered on and the amount of memory must be reduced.

This version of the action has the Power Off Allowed flag set to true. You can select this Power Off Allowed version of the action when you create or edit alerts and associate the alert with a recommendation. When the Power Off Allowed version of this action is automated, you do not select this version of the action.

If Hot Plug is activated on the virtual machines, then power off is not required. If power off is required and VMware Tools is installed, then the virtual machines are shut down before they are powered off.

**Where You Run the Action**

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click **Views**.
- From the left menu click **Operations › Configurations**, and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

**Action Options**

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.<br><br>If you modify a value, the check box is selected. The check box must be selected to activate the OK button. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| New CPU | Number of CPUs when the action is completed. If the value is less than 1 or a value not supported for the virtual machine in vCenter, and the virtual machine is powered on and Hot Add is not activated, the number of CPUs does not change and Recent Tasks shows the action as failed. If the virtual machine is powered off when you submit an unsupported value, the task reports success, but the virtual machine will fail when you run a power on action.<br><br>The value that appears is the calculated suggested size. If the target virtual machine is new or offline, this value is the current number of CPUs. If VMware Aria OperationsV Mware Cloud Foundation Operations has been monitoring the virtual machine for six or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric. |
| Current CPU | Number of configured CPUs. |
| Power State | Indicates whether the virtual machine is powered on or powered off. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Power Off Allowed | If selected, the action shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without regard for the state of the operating system.<br><br>In addition to whether the action shuts down or powers off a virtual machine, you must consider whether the object is powered on and what settings are applied.<br><br>See Working with Actions That Use Power Off section in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* . |
| Snapshot | Creates a snapshot before changing the number of CPUs. Use this option if you need a snapshot to which you can revert the virtual machine if the action does not produce the expected results.<br><br>The name of the snapshot is supplied in the Recent Tasks messages for the action.<br><br>If the CPU is changed with CPU Hot Plug activated, then the snapshot is taken with the virtual machine is running, which consumes more disk space. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 392: Task ID Dialog Box**

| Option | Description |
|---|---|
| OK | To close the dialog box without further action, click **OK**. |
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |

### Set Memory Resources for Virtual Machine Action

The Set Memory Resources for VM action is used to modify the memory reservation and memory limit on virtual machines. You modify the memory reservation and limit to manage resources in your environment, either to reclaim unused resources or to ensure that your virtual machines have the resources they need to run efficiently.

## How the Action Works

The Set Memory Resources for VM action determines how memory resources are allocated to the virtual machine. The reservation value is the minimum amount of guaranteed memory allocated for the virtual machine. The limit is the maximum amount of memory that the virtual machine can consume.

The reservation and limit values in vCenter are set in megabytes. VMware Aria OperationsVMware Cloud Foundation Operations calculates and reports on memory in kilobytes. When you run this action, the values are presented in kilobytes so that you can implement recommendations from VMware Aria OperationsVMware Cloud Foundation Operations.

To run the action, all options must be configured in the dialog box for the objects on which your are running the action. If you are changing one option to a new value, but not another option, ensure that the option that you do not want to change is configured with the current value.

## Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click **Views**.
- From the left menu click **Operations › Configurations** and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

## Action Options

To ensure that you are submitting the action for the right objects, review the following information.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.<br><br>If you modify a value, the check box is selected. The check box must be selected to activate the OK button. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| New Resv (KB) | Amount of memory in kilobytes reserved for the virtual machine when the action is finished. The new reservation value must be less than or equal to the new limit value unless your new limit is unlimited (-1).<br><br>The reservation supports the following possible values:<br><br>• If you set the value to 0, the virtual machine is allocated only the currently configured amount of RAM.<br>• If you add or remove reserved memory, the value must be evenly divisible by 1024. |
| Current Resv (KB) | Amount of memory in kilobytes that is configured as the guaranteed memory for the virtual machine. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| New Limit (KB) | Maximum amount of memory in kilobytes that the virtual machine can consume when the action is completed.<br><br>The limit supports the following possible values:<br><br>• If you set the value to 0, then the maximum memory is no greater than the allocated reservation amount.<br>• If you set the value to -1, then the virtual machine memory is unlimited.<br>• It you increase or decrease the limit, the value must be evenly divisible by 1024. |
| Current Limit (KB) | Maximum amount of memory that the virtual machine is currently allowed to consume. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 393: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Set CPU Count for Virtual Machine Action

The Set CPU action modifies the number of vCPUs on a virtual machine. You increase the number of CPUs to address performance problems or decrease the number of CPU to reclaim resources.

#### How the Action Works

The Set CPU Count action shuts down or powers off the target virtual machines. If you are decreasing the CPU count, the action is required. This action creates a snapshot if you request it, changes the number of vCPUs based on the new CPU count you provided, and returns the virtual machines to their original power states.

An alternative form of the Set CPU Count for Virtual Machine action is available for automation. This action can run when the virtual machine is powered on or off.

Use this version of the action if the automated action has permission to power off the virtual machine, and hot add of memory is not activated on the virtual machine. With hot add activated, you can add CPUs, but you cannot remove them.

This version of the action is required if a virtual machine is powered on and the number of CPUs must be reduced.

This version of the action has the Power Off Allowed flag set to true. You can select this Power Off Allowed version of the action when you create or edit alerts and associate the alert with a recommendation. When the Power Off Allowed version of this action is automated, you do not select this version of the action.

If Hot Plug is activated on the virtual machines, then power off is not required. If power off is required and VMware Tools are installed, then the virtual machines are shut down before they are powered off.

**Where You Run the Action**

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click **Views**.
- From the left menu click **Operations › Configurations**, and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

**Action Options**

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.<br><br>If you modify a value, the check box is selected. The check box must be selected to activate the OK button. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| New CPU | Number of CPUs when the action is completed. If the value is less than 1 or a value not supported for the virtual machine in vCenter, and the virtual machine is powered on and Hot Add is not activated, the number of CPUs does not change and Recent Tasks shows the action as failed. If the virtual machine is powered off when you submit an unsupported value, the task reports success, but the virtual machine will fail when you run a power on action.<br><br>The value that appears is the calculated suggested size. If the target virtual machine is new or offline, this value is the current number of CPUs. If VMware Aria OperationsV Mware Cloud Foundation Operations has been monitoring the virtual machine for six or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric. |
| Current CPU | Number of configured CPUs. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Power State | Indicates whether the virtual machine is powered on or powered off. |
| Power Off Allowed | If selected, the action shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without regard for the state of the operating system.<br><br>In addition to whether the action shuts down or powers off a virtual machine, you must consider whether the object is powered on and what settings are applied.<br><br>See Working with Actions That Use Power Off section in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* . |
| Snapshot | Creates a snapshot before changing the number of CPUs. Use this option if you need a snapshot to which you can revert the virtual machine if the action does not produce the expected results.<br><br>The name of the snapshot is supplied in the Recent Tasks messages for the action.<br><br>If the CPU is changed with CPU Hot Plug activated, then the snapshot is taken with the virtual machine is running, which consumes more disk space. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 394: Task ID Dialog Box**

| Option | Description |
|---|---|
| OK | To close the dialog box without further action, click **OK**. |
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |

### Set CPU Resources for Virtual Machine Action

The Set CPU Resources for VM action is used to modify the CPU reservation and CPU limit on virtual machines. You modify the CPU reservation and limit to manage workload demands in your environment.

**How the Action Works**

The Set CPU Resources for VM action determines how CPU resources can be allocated to the virtual machines. The reservation limit is the minimum amount of guaranteed CPU resources allocated to the virtual machine. The limit is the maximum amount of CPU resources that the virtual machine can consume.

To run the action, all options where you configure a value must contain a value for the objects that you want to change. If you are changing one option to a new value, but not another option, ensure that the option that you are not changing is configure with the current value.

**Where You Run the Action**

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and click **Views**.
- From the left menu click **Operations › Configuration**, and then click the **Inventory Management** tab. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

**Action Options**

To ensure that you are submitting the action for the right objects, review the following information.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.<br><br>If you modify a value, the check box is selected. The check box must be selected to activate the OK button. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| New Resv (MHz) | Amount of CPU resources in megahertz reserved for the virtual machine when the action is finished. The new reservation value must be less than or equal to the new limit value unless your new limit is unlimited (-1).<br><br>The reservation supports the following possible values:<br><br>• If you set the value to 0, the virtual machine is allocated only the configured CPU consumption level.<br>• If you add or removed reserved CPU consumption, supply a positive integer unless you set the value to 0. |
| Current Resv (MHz) | Amount of CPU resources that is configured as the guaranteed CPU resources for the virtual machine. |
| New Limit (MHz) | Maximum amount of CPU consumption in megahertz that the virtual machine can consume when the action is completed. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | The limit supports the following possible values:<br>• If you set the value to 0, the maximum CPU consumption is not greater than the allocated reservation amount.<br>• If you set the value to -1, then the virtual machine CPU consumption is unlimited.<br>• If you add or remove CPU consumption limits, supply a positive integer, unless you set the value to 0 or -1. |
| Current Limit (MHz) | Maximum amount of CPU that the virtual machine can consume. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 395: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

### Set CPU Count and Memory for Virtual Machine Action

The Set CPU Count and Memory for VM action is used to add or remove CPUs and memory on virtual machines with only one power off of the virtual machines to perform the combined actions. You modify the CPU and memory to address performance problems or to reclaim resources.

#### How the Action Works

The Set CPU Count and Memory action powers off the target virtual machines. The action also creates a snapshot when requested and changes the number of vCPUs and memory based on the new CPU count and memory values you provided. As well, the action returns the virtual machines their original power states.

An alternative form of the Set CPU Count and Memory for Virtual Machine action is available for automation. This version of the action has the Power Off Allowed flag set to true so that the action is available for automation and can run when the virtual machine is in the powered on state. You can select the Power Off Allowed version of the action when you create or edit alerts and associate the alert with a recommendation. When the Power Off Allowed version of this action is automated, you do not select this version of the action.

If Hot Plug is activated on the virtual machines, then power off is not required. If power off is required and VMware Tools are installed, then the virtual machines are shut down before they are powered off.

To run the action, all options where you configure a value must contain a value for the objects that you want to change. If you are changing one option to a new value, but not another option, ensure that the option that you are not changing is configure with the current value.

**Where You Run the Action**

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and click **Views**.
- From the left menu click **Operations › Configurations**, and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

**Action Options**

Review the following information about the virtual machines to ensure that you are submitting the action for the correct objects.

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected.<br><br>If you modify a value, the check box is selected. The check box must be selected to activate the OK button. |
| Name | Name of the virtual machine as it appears in the environment inventory. |
| New CPU | Number of CPUs when the action is completed. If the value is less than 1 or a value not supported for the virtual machine in vCenter, and the virtual machine is powered on and Hot Add is not activated, the number of CPUs does not change and Recent Tasks shows the action as failed. If the virtual machine is powered off when you submit an unsupported value, the task reports success, but the virtual machine will fail when you run a power on action.<br><br>The value that appears is the calculated suggested size. If the target virtual machine is new or offline, this value is the current number of CPUs. If VMware Aria OperationsV Mware Cloud Foundation Operations has been monitoring the virtual machine for six or more hours, depending on your environment, the value that appears is the CPU Recommended Size metric. |
| Current CPU | Number of configured CPUs. |
| Power State | Indicates whether the virtual machine is powered on or powered off. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Power Off Allowed | If selected, the action shuts down or powers off the virtual machine before modifying the value. If VMware Tools is installed and running, the virtual machine is shut down. If VMware Tools is not installed or not running, the virtual machine is powered off without regard for the state of the operating system.<br><br>In addition to whether the action shuts down or powers off a virtual machine, you must consider whether the object is powered on and what settings are applied.<br><br>See Working with Actions That Use Power Off section in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* . |
| Snapshot | Creates a snapshot before changing the number of CPUs. Use this option if you need a snapshot to which you can revert the virtual machine if the action does not produce the expected results.<br><br>The name of the snapshot is supplied in the Recent Tasks messages for the action.<br><br>If the CPU is changed with CPU Hot Plug activated, then the snapshot is taken with the virtual machine is running, which consumes more disk space. |
| Host | Name of the host on which the virtual machine is running. |
| Adapter Instance | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 396: Task ID Dialog Box**

| Option | Description |
|---|---|
| OK | To close the dialog box without further action, click **OK**. |
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |

### Delete Unused Snapshots for Virtual Machine Action

The Delete Unused Snapshots for Virtual Machines action in VMware Aria OperationsVMware Cloud Foundation Operations deletes snapshots that are older than the specified age from your datastores. Deleting unused snapshots reclaims wasted space in your environment.

## How the Action Works

The Delete Unused Snapshots for Virtual Machine action comprises two dialog boxes. The first dialog box allows you to select the snapshot age criteria, which must be greater than one day. The second step allows you to select the snapshots to delete, and runs the Delete Unused Snapshots for Virtual Machine action.

The number of days that you specify for each virtual machine is the age of the snapshots based on the creation date. The Delete Unused Snapshots for Virtual Machine action retrieves the snapshot and displays the snapshot name, space consumed, and location so that you can evaluate the snapshots before you delete them.

When you click **Begin Action**, VMware Aria OperationsVMware Cloud Foundation Operations displays a dialog box to indicate that the action has started. To track the status of the action, click the link in the dialog box and view the state of the action in **Recent Tasks**.

## Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click the **Views** tab.
- From the left menu click **Inventory**, select an object, click the **Details** tab, and then click the **List** tab.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

## Action Options

To ensure that you are submitting the action for the right objects, review the following information.

You first retrieve snapshots based on age, then select the snapshots to delete.

**Table 397: Retrieve Snapshots**

| Option | Description |
|---|---|
| Name | Name of the virtual machine on which you are running the Delete Unused Snapshots for VM action. |
| Days Old | Age of the snapshots to be deleted. This action retrieves snapshots for the virtual machine that are older than one day. |
| Host | Name of the host with which the virtual machine is associated. |
| Parent vCenter | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

Select the snapshots to delete.

**Table 398: Delete Snapshots**

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. |
| VM Name | Name of the virtual machine from which the snapshot was created. |
| Snapshot Name | Name of the snapshot in the datastore. |
| Snapshot Space (MB) | Number of megabytes consumed by the snapshot. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| Snapshot Create Time | Date and time when the snapshot was created. |
| Snapshot Age | Age of the snapshot in days. |
| Datacenter Name | Name of the data center with which the datastore is associated. |
| Datastore Name | Name of the datastore where the snapshot is managed. |
| Host Name | Name of the host with which the datastore is associated. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 399: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

The Delete Unused Snapshots action creates a job for the retrieve snapshots action, and a job for the delete snapshots action.

### Delete Unused Snapshots for Datastore Action

The Delete Unused Snapshots for Datastore action in VMware Aria OperationsVMware Cloud Foundation Operations deletes snapshots that are older than the specified age from your datastores. Deleting unused snapshots reclaims wasted space in your environment.

#### How the Action Works

The Delete Unused Snapshots for Datastore action comprises two dialog boxes. The first dialog box allows you to select the snapshot age criteria, which must be greater than one day. The second step allows you to select the snapshots to delete, and runs the Delete Unused Snapshots for Datastore action.

The number of days that you specify for each datastore is the age of the snapshots based on the creation date. The Delete Unused Snapshots dialog box provides details regarding snapshot name, space consumed, and location so that you can evaluate the snapshots before you delete them.

When you click **Begin Action**, VMware Aria OperationsVMware Cloud Foundation Operations displays a dialog box to indicate that the action has started. To track the status of the action, click the link in the dialog box and view the state of the action in **Recent Tasks**.

#### Where You Run the Action

For the supported objects and object levels, this action is available in the following locations in VMware Aria OperationsVMware Cloud Foundation Operations:

- Embedded just below the top menu.
- From the left menu, click **Inventory**, select an object, click the **Details** tab, and then click the **List** tab.
- From the left menu click **Operations › Configurations**, and then click the **Inventory Management** tile. Select an object in the list.
- In configured alert recommendations.
- In the Object List and Topology Graph dashboard widgets.

**Action Options**

To ensure that you are submitting the action for the right objects, review the following information.

You first retrieve snapshots based on age, then select the snapshots to delete.

**Table 400: Retrieve Snapshots**

| Option | Description |
|---|---|
| Name | Name of the datastore on which you are running the delete snapshot action. |
| Days Old | Age of the snapshots to be deleted. This action retrieves snapshots for the datastore that are older than one day. |
| Host | Name of the host with which the datastore is associated. |
| Parent vCenter | Name of the VMware Adapter as it is configured in VMware Aria OperationsVMware Cloud Foundation Operations. The adapter manages the communication with the vCenter instance. |

Select the snapshots to delete.

**Table 401: Delete Snapshots**

| Option | Description |
|---|---|
| Selected objects | Check box indicates whether the action is applied to the object. To not run the action on one or more objects, deselect the associated check boxes. This option is available when two or more objects are selected. |
| Datastore Name | Name of the datastore where the snapshot is managed. |
| Snapshot Name | Name of the snapshot in the datastore. |
| Snapshot Space (MB) | Number of megabytes consumed by the snapshot. |
| Snapshot Create Time | Date and time when the snapshot was created. |
| Snapshot Age | Age of the snapshot in days. |
| Datacenter Name | Name of the data center with which the datastore is associated. |
| Host Name | Name of the host with which the datastore is associated. |
| VM Name | Name of the virtual machine from which the snapshot was created. |

After you click **OK**, the next dialog box provides the task ID and a link to the task list.

**Table 402: Task ID Dialog Box**

| Option | Description |
|---|---|
| Recent Tasks | To view the status of the job and verify that the job finished, click **Recent Tasks**. |
| OK | To close the dialog box without further action, click **OK**. |

The Delete Unused Snapshots action creates a job for the retrieve snapshots action, and a job for the delete snapshots action.

**Power On, Power Off, and Reboot Actions**

You can run Power On, Power Off, and Reboot actions for AWS, Azure, and GCP using VMware Aria OperationsVMware Cloud Foundation Operations. These actions allow you to manage your objects and respond quickly to alerts. You can power on an instance to start collecting data, run applications, or verify the actions that were run on already powered-down machines. You can power off or reboot an instance to manage resources and reclaim wasted space.

- To run actions for AWS, ensure your account ID has the following privileges: **ec2:StartInstances**, **ec2:StopInstances** , and **ec2:RebootInstances**.
- To run actions for Azure, ensure you have one of the following privileges: **contributor**, **owner**, **virtual machine contributor**, and **avere contributor**.
- To run actions for GCP, ensure you have one of the following privileges: **owner**, **compute admin**, and **Compute Instance Admin (v1)**.

1. To perform actions on a single instance:
   a) From the left menu, click **Administration › Integrations**.
   b) In the **Accounts** tab, click the vertical ellipses against the AWS, Microsoft Azure, or GCP Cloud Account, and then select **Object Details**.
   The Object Browser page is displayed.
   c) Locate the EC2 Instance, Azure Virtual Machine, or CE Instance and click **Actions**.
2. To perform actions on multiple instances at once:
   a) From the left menu click **Inventory**.
   b) Navigate to the AWS, Microsoft Azure, or GCP instance and filter by object type.
   c) Select multiple objects and click the **Actions** icon.
3. Select one of the following actions.

   The Actions vary based on the instance you are using.

   **NOTE**
   To start an instance that is in a powered-off state, use the Power On action. The Power Off and Reboot actions appear only when the instance is in the Power On state.

| Accounts | Actions |
|---|---|
| AWS: EC2 Instance | • Power On Instance<br>• Power Off Instance<br>• Reboot Instance<br>**NOTE**<br>These actions are available only if you activate actions while configuring an AWS account. For more information, see the "Add a Cloud Account for AWS" topic, in the *VMware Aria Operations Configuration Guide*. |
| Microsoft Azure: Azure Virtual Machine | • Power On VM<br>• Power Off VM<br>• Reboot VM |

*Table continued on next page*

*Continued from previous page*

| Accounts | Actions |
|---|---|
| | **NOTE**<br>These actions are available only if you activate actions while configuring a Microsoft Azure account. For more information, see the "Add a Cloud Account for Microsoft Azure" topic, in the *VMware Aria Operations Configuration Guide*. |
| GCP: CE Instance | • Power On CE Instance<br>• Power Off CE Instance<br>• Reboot CE Instance<br>**NOTE**<br>These actions are available only if you activate actions while configuring a GCP account. For more information, see the "Configuring Google Cloud Platform" topic, in the *VMware Aria Operations Configuration Guide*. |

4. Click **Begin Action** and then, click **OK**.

Click **Recent Tasks** under **Administration › Control Panel** to verify the status of the action.

### Troubleshoot Actions in VMware Aria OperationsVMware Cloud Foundation Operations

If you are missing data or cannot run actions from VMware Aria OperationsVMware Cloud Foundation Operations, review the troubleshooting options.

Verify that your vCenter Adapter is configured to connect to the correct vCenter instance, and configured to run actions. See Configure a vCenter Server Cloud Account section in *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.*
Verify that your vCenter Adapter is configured to connect to the correct vCenter instance, and configured to run actions. See *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide.* .

### Actions Do Not Appear on Object

An action might not appear on an object, such as a host or virtual machine, because VMware Aria Automation is managing that object.

Actions such as Rebalance Container might not appear in the drop-down menu when you view the actions for your data center.

- If a data center is managed by VMware Aria Automation, actions do not appear.
- If a data center is not managed by VMware Aria Automation, you can act on the virtual machines that VMware Aria Automation is not managing.

When VMware Aria Automation manages the child objects of a data center or custom data center container, the actions that are normally available on those objects do not appear. They are not available because the action framework excludes actions on objects that VMware Aria Automation manages. You cannot turn on or turn off the exclusion of actions on objects that VMware Aria Automation manages. This behavior is normal.

If you removed the VMware Aria Automation adapter instance, but did not select the **Remove related objects** check box, the actions are still disabled.

Make actions available on the objects in your data center or custom data center in one of two ways. Either confirm that VMware Aria Automation is not managing the objects, or perform the steps in this procedure to remove the VMware Aria Automation adapter instance.

1. To allow actions on an object, go to your VMware Aria Automation instance.
2. Perform the action in VMware Aria Automation, such as to move a virtual machine.

### Missing Column Data in Actions Dialog Boxes

Data is missing for one or more objects in an Actions dialog box, making it difficult to determine if you want to run the action.

When you run an action one or more objects, some of the fields are empty.
There are two possible causes: 1) the VMware vSphere adapter has not collected the data from the vCenter instance that manages the object. 2) the current VMware Aria OperationsVMware Cloud Foundation Operations user does not have privileges to view the collected data for the object.

1. Verify that VMware Aria OperationsVMware Cloud Foundation Operations is configured to collect the data.
2. Verify that you have the privileges necessary to view the data.

### Missing Column Data in the Set Memory for VM Dialog Box

The read-only data columns do not display the current values, which makes it difficult to specify properly a new memory value.

Current (MB) and Power State columns do not display the current values, which are collected for the managed object.
The adapter responsible for collecting data from the vCenter on which the target virtual machine is running has not run a collection cycle and collected the data. This omission can occur when you recently created an VMware adapter instance for the target vCenter and initiated an action. The VMware vSphere adapter has a five-minute collection cycle.

1. After you create a VMware adapter instance, wait an extra five minutes.
2. Rerun the **Set Memory for VM** action.
   The current memory value and the current power state appear in the dialog box.

### Host Name Does Not Appear in Action Dialog Box

When you run an action on a virtual machine, the host name is blank in the action dialog box.

When you select virtual machine on which to run an action, and click the **Action** button, the dialog box appears, but the Host column is empty.
Although your user role is configured to run action on the virtual machines, you do not have a user roll that provides you with access to the host. You can see the virtual machines and run actions on them, but you cannot see the host data for the virtual machines. VMware Aria OperationsVMware Cloud Foundation Operations cannot retrieve data that you do not have permission to access.
You can run the action, but you cannot see the host name in the action dialog boxes.

### Monitor Recent Task Status

The Recent Task status includes all the tasks initiated from VMware Aria OperationsVMware Cloud Foundation Operations. You use the task status information to verify that your tasks finished successfully or to determine the current state of tasks.

You ran at least one action as part of an alert recommendation or from one of the toolbars. See Run Actions from Toolbars in VMware Aria OperationsVMware Cloud Foundation Operations.

You can monitor the status of tasks that are started when you run actions, and investigate whether a task finished successfully.

1. From the left menu, click **Administration › Control Panel**, and then select the **Recent Tasks** tile.
2. To determine if you have tasks that are not finished, click the **Status** column and sort the results.

| Option | Description |
|---|---|
| `In Progress` | Indicates running tasks. |
| `Completed` | Indicates finished tasks. |
| `Failed` | Indicates incomplete tasks on at least one object when started on multiple objects. |
| `Maximum Time Reached` | Indicates timed out tasks. |

3. To evaluate a task process, select the task in the list and review the information in the **Details of Task Selected** pane.
   The details appear in the Messages pane. If the information message includes `No action taken`, the task finished because the object was already in the requested state.
4. To view the messages for an object when the task included several objects, select the object in the Associated Objects list.

   To clear the object selection so that you can view all the messages, press the space bar.

Troubleshoot tasks with a status of `Maximum Time Reached` or `Failed` to determine why a task did not run successfully. See Troubleshoot Failed Tasks.

### Recent Tasks in VMware Aria OperationsVMware Cloud Foundation Operations

The status of the tasks that were recently initiated from VMware Cloud Foundation OperationsVMware Cloud Foundation Operationsappears in the Recent Task list. You can determine whether a task is finished, still in process, or failed.

#### How Recent Tasks Work

The Recent Tasks page reports on logged task events, and the log entries appear in the messages area so that you can troubleshoot failed tasks.

#### Where You View Recent Tasks

From the left menu, select **Administration**, then click **Recent Tasks** under **Control Panel**.

#### Recent Task Options

Review the information in the task list to determine if a task is completed or if you must troubleshoot a failed task. To see the details about a task, select the task in the list and review the associated objects and task messages.

**Table 403: Task List**

| Option | Description |
|---|---|
| Export | Exports the selected task to an XML file.<br><br>The exported information, which includes the messages, is useful when you are troubleshooting a problem. |
| Edit Properties | Determines how long the recent task data is retained in your system.<br><br>Set the number of days that VMware Aria OperationsVMware Cloud Foundation Operations keeps the data, after which it is purged from the system. The default value is 90 days. |
| Status drop-down menu | Filters the list based on the status value. |
| All Filters | Filters the list based on the selected column and the provided values. You can filter the tasks based on the following criteria.<br>• Task<br>• Started Time<br>• Completed Time<br>• Automated<br>• Object Name<br>• Object Type<br>• Event Source<br>• Source Type<br>• Submitted By<br>• Task ID |
| Filter (Object Name) | Limits the tasks in the list to those that match the entered string.<br><br>The search is based on a partial entry. For example, if you enter `vm`, objects such as `vm001` and `acctvm_east` are included. |
| Task | Name of the task.<br><br>For example, Set CPU Count for VM. |
| Status | State of the task.<br><br>Possible states include the following values:<br>• Completed. Task completed successfully on the target objects.<br>• In Progress. Task is running on the target objects.<br>• Failed. Task failed to run on the target objects. If the task started, the reasons for failure might include a faulty script, a script timed out, or actions are not taken. If the task did not start and immediately reports as failed, the reasons might include that the task was not able to start or the script was not found. If the task was not initiated on the target object, it might have failed because of communication or authentication errors. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | • Maximum Time Reached. Task is running past the amount of time that is the default or configured value. To determine the status, you must troubleshoot the initiated action.<br>• Not Dispatched. The action adapter was not found.<br>• Started. Task is initiated on the object.<br>• Unknown. An error occurred while running the action, but the error was not captured in the task logs. To investigate this status further, check the VMware Aria OperationsVMware Cloud Foundation Operations support logs for the vCenter Adapter, available in the **Administration** area, and check the target system. |
| Started Time | Date and time when the task started. |
| Completed Time | Date and time when the task finished.<br><br>A completed date does not appear if the task failed or if the maximum timeout is reached. |
| Automated | Indicates whether the action in the task list is automated, indicated by `Yes` or `No`. |
| Object Name | Object on which the task was started. |
| Object Type | Type of object on which the task was started. |
| Event Source | The UUID or the name of the event that triggered the action automatically. when an event is triggered that is associated to the recommendation, it triggers the action without the user intervention.<br><br>For example, you can automate Alert recommendations that have an associated action. Automation is disabled by default. You configure automation in the **Override Alert / Symptom Definitions** area of a policy when you create or edit the policy in **Operations › Configurations**, and then click the **Policy Definition** tile.<br><br>An administrator who has the **Automation** role has permission to automate actions in the **Override Alert / Symptom Definitions** area of the policy workspace. |
| Source Type | Authentication source that the user who started the task used when accessing VMware Aria OperationsVMware Cloud Foundation Operations. |
| Submitted By | Name of the user who initiated the task. This column displays the automationAdmin user account for automated actions that are triggered by alerts. |
| Task ID | ID generated when the task, which included one or more actions, was started. |

*Table continued on next page*

*Continued from previous page*

| Option | Description |
|---|---|
| | The task ID is unique for the task for each adapter. If a task includes tasks that ran using two adapters, you see two task IDs.<br><br>If the task is a delete snapshot action, two task IDs are generated. One ID is for the retrieve snapshots based on date task, and the other ID is for the delete selected snapshots task. |

The Associated Objects are the objects on which the selected task ran.

**Table 404: Associated Objects for Selected Task Details**

| Option | Description |
|---|---|
| Object Name | Detailed list of objects that are included in the task selected in the task list.<br><br>If the task ran on only one object, the list includes one object. If the task ran on multiple objects, each object is listed on a separate row. |
| Object Type | Type of object for each object name. |
| Status | Current state of the task. |
| The following are columns displayed for a Workload Optimization task | |
| Source Cluster | The source Cluster from which the VM is moved during Workload Optimization. |
| Destination Cluster | The destination Cluster to which the VM is moved during Workload Optimization. |
| Source Host | The source Host from which the VM is moved during Workload Optimization. |
| Destination Host | The destination Host to which the VM is moved during Workload Optimization. |
| Source Datastore | The source Datastore from which the VM is moved during Workload Optimization. |
| Destination Datastore | The destination Datastore to which the VM is moved during Workload Optimization. |
| Completion Date | The date when the task was completed. |

The Messages are the log of the task as it ran. If the task does not finish successfully, use the logs to identify problems.

**Table 405: Messages for Selected Task Details**

| | |
|---|---|
| Severity drop-down menu | Limits the messages based on the Severity value. |
| Filter (Message) | Limits the message in the list to those that match the entered string. |

*Table continued on next page*

*Continued from previous page*

| | |
|---|---|
| | The search is based on a partial entry. For example, if you enter `id`, then messages that contain `Task ID` and the phrase `did not complete` are included. |
| Severity | Message level in the logs.<br><br>The severity includes the following values:<br>• All. Displays all the messages.<br>• Error. Messages generated during a task failure.<br>• Warning. Messages generated as warning when the task is in progress.<br>• Information. Messages added to logs as the task is processed. |
| Time | Date and time the entry was added to the log. |
| Message | Text of the log entry.<br><br>Use the information in the message to determine why a task failed, and to begin to troubleshoot and resolve the failure.<br><br>The messages appear with the most recent entry at the top of the list if you do not sort the columns. |

**Troubleshoot Failed Tasks**

If tasks fail to run in VMware Aria OperationsVMware Cloud Foundation Operations, review the Recent Tasks page and troubleshoot the task to determine why it failed.

This information is a general procedure for using the information in Recent Tasks to troubleshoot problems identified in the tasks.

**Determine If a Recent Task Failed**

The Recent Tasks provide the status of action tasks initiated from VMware Aria OperationsVMware Cloud Foundation Operations. If you do not see the expected results, review the tasks to determine if your task failed.

1. From the left menu, click **Administration › Control Panel**, then click the **Recent Tasks** tab.
2. Select the failed task in the task list.
3. In the Messages list, locate the occurrences of `Script Return Result: Failure` and review the information between this value and `<-- Executing:[script name] on {object type}.`

   `Script Return Result` is the end of action run and `<-- Executing` indicates the beginning. The information provided includes the parameters that are passed, the target object, and unexpected exceptions that you can use to identify the problem.

**Troubleshooting Maximum Time Reached Task Status**

An action task has a `Maximum Time Reached` status and you do not know the status of the task.

The Recent Tasks list indicates that a task had a status of `Maximum Time Reached`.

The task is running past the amount of time that is the default or configured value. To determine the latest status, you must troubleshoot the initiated action.

The task is running past the amount of time that is the default or configured value for one of the following reasons:

- The action is exceptionally long running and did not finish before the threshold timeout was reached.
- The action adapter did not receive a response from the target system before reaching the timeout. The action might have completed successfully, but the completion status was not returned to VMware Aria OperationsVMware Cloud Foundation Operations.
- The action did not start correctly.
- The action adapter might have an error and be unable to report the status.

To determine whether the action completed successfully, check the state of the target object. If it did not complete, continue investigating to find the root cause.

### Troubleshooting Set CPU or Set Memory Failed Tasks

An action task for Set CPU Count or Set Memory for VM has a `Failed` status in the recent task list because power off is not allowed.

The Recent Tasks list indicates that a Set CPU Count, Set Memory, or Set CPU and Memory task has a status of `Failed`. When you evaluate the Messages list for the selected task, you see this message.

```
Unable to perform action. Virtual Machine found

     powered on, power off not allowed.
```

When you increase the memory or CPU count, you see this message.

```
Virtual Machine found powered on, power off not allowed, if hot add is

     enabled the hotPlugLimit is exceeded.
```

You submitted the action to increase or decrease the CPU or memory value without selecting the **Allow Power Off** option. When you ran the action where a target object is powered on and where **Memory Hot Plug** is not activated for the target object in vCenter, the action fails.

1. Either activate **Memory Hot Plug** on your target virtual machines in vCenter or select **Allow Power Off** when you run the Set CPU Count, Set Memory, or Set CPU and Memory actions.
2. Check your hot plug limit in vCenter.

### Troubleshooting Set CPU Count or Set Memory with Powered Off Allowed

A Set CPU Count, Set Memory, or a Set CPU Count and Set Memory action indicates that the action failed in Recent Tasks.

When you run an action that changes the CPU count, the memory, or both, the action fails. It fails even though Power Off Allowed was selected, the virtual machine is running, and the VMware Tools are installed and running.
The virtual machine must shut down the guest operating system before it powers off the virtual machine to make the requested changes. The shutdown process waits 120 seconds for a response from the target virtual machine, and fails without changing the virtual machine.

1. To determine if it has jobs running that are delaying the implementation of the action, check the target virtual machine in vCenter.
2. Retry the action from VMware Aria OperationsVMware Cloud Foundation Operations.

### Troubleshooting Set CPU Count and Memory When Values Not Supported

If you run the Set CPU Count or Set Memory actions with an unsupported value on a virtual machine, the virtual machine might be left in an unusable state. That outcome requires you to resolve the problem in vCenter.

You cannot power on a virtual machine after you successfully run the Set CPU Count or Set Memory actions. When you review the messages in Recent Tasks for the failed Power On VM action, you see messages stating that the host does not support the new CPU count or new memory value.
Because of the way that vCenter validates changes in the CPU and memory values, you can use the VMware Aria OperationsVMware Cloud Foundation Operations actions to change the value to an unsupported amount. This change can happen when you run the action when the virtual machine is powered off.

If the object was powered on, the task fails, but rolls back any value changes and powers the machine back on. If the object was powered off, the task succeeds and the value is changed in vCenter. However, the target object is left in a state where you cannot power it on using either actions or the vCenter without manually changing the CPU or memory to a supported value.

1. From the left menu, click **Administration › Control Panel**, and then select the **Recent Tasks** tile.
2. In the task list, locate your failed Power On VM action, and review the messages associated with the task.
3. Look for a message that indicates why the task failed.
   For example, if you ran a Set CPU Count action on a powered off virtual machine to increase the CPU count from 2 to 4, but the host does not support 4 CPUs. The Set CPU tasks reported that it completed successfully in recent tasks. However, when you attempt to power on the virtual machine, the tasks fails. In this example, the message is `Virtual machine requires 4 CPUs to operate, but the host hardware only provides 2.`
4. Click the object name in the Recent Task list.
   The main pane updates to display the object details for the selected object.
5. Click the **Actions** menu on the toolbar and click **Open Virtual Machine in vSphere Client**.
   The vSphere Web Client opens with the virtual machine as the current object.
6. In the vSphere Web Client, click the **Manage** tab and click **VM Hardware**.
7. Click **Edit**.
8. In the Edit Settings dialog box, change the CPU count or memory to a supported value and click **OK**.
   You can now power on the virtual machine from the Web client or from VMware Aria OperationsVMware Cloud Foundation Operations.

### Troubleshooting Set CPU Resources or Set Memory Resources When the Value Is Not Supported

If you run the Set CPU Resources action with an unsupported value on a virtual machine, the task fails and an error appears in the Recent Task messages.

The Recent Tasks list indicates that a Set CPU Resource or Set Memory Resource action has a state of `Failed`. When you evaluate the Messages list for the selected task, you see a message similar to the following examples.

```
RuntimeFault exception, message:[A specified parameter was not correct.
spec.cpuAllocation.reservation]
```

```
RuntimeFault exception, message:[A specified parameter was not correct.
spec.cpuAllocation.limits]
```

You submitted the action to increase or decrease the CPU or memory reservation or limit value with an unsupported value. For example, if you supplied a negative integer other than -1, which sets the value to unlimited, vCenter cannot make the change and the action failed.

1. Run the action with a supported value.

   The supported values for reservation include 0 or a value greater than 0. The supported values for limit include -1, 0, or a value greater than 0.

### Troubleshooting Set CPU Resources or Set Memory Resources When the Value Is Too High

You run the Set CPU Resources or Set Memory Resources action and the task fails with an error appearing in the Recent Tasks messages. The reason might be that you entered a value that is greater than the value that your vCenter instance supports.

The Recent Tasks list indicates that a Set CPU Resource or Set Memory Resource action has a state of `Failed`. When you evaluate the Messages list for the selected task, you see messages similar to the following examples.

If you are working with Set CPU Resources, the information message is similar to the following example, where 1000000000 is the supplied reservation value.

```
Reconfiguring the Virtual Machine Reservation to:[1000000000] Mhz
```

The error message for this action is similar to this example.

```
RuntimeFault exception, message:[A specified parameter was not correct: reservation]
```

If you are working with Set Memory Resources, the information message is similar to the following example, where 1000000000 is the supplied reservation value.

```
Reconfiguring the Virtual Machine Reservation to:[1000000000] (MB)
```

The error message for this action is similar to this example.

```
RuntimeFault exception, message:[A specified parameter was not correct.
spec.memoryAllocation.reservation]
```

You submitted the action to change the CPU or memory reservation or limit value to a value greater than the value supported by vCenter, or the submitted reservation value is greater than the limit.

1. Run the action using a lower value.

### Troubleshooting Set Memory Resources When the Value Is Not Evenly Divisible by 1024

If you run the Set Memory Resources action with a value that cannot convert from kilobytes to megabytes, the task fails and an error appears in the Recent Task messages.

The Recent Tasks list indicates that a Set Memory Resource action has a state of `Failed`. When you evaluate the Messages list for the selected task, you see a message similar to the following example.

```
Parameter validation;[newLimitKB] failed conversion to (MB, (KB)[2000] not evenly
divisible by 1024.
```

Because vCenter manages memory reservations and limit values in megabytes, but VMware Aria OperationsVMware Cloud Foundation Operations calculates and reports on memory in kilobytes, you must provide a value in kilobytes that is directly convertible to megabytes. To do that, the value must be evenly divisible by 1024.

1. Run the action where the reservation and limit values are configured with supported values.

   The supported values for reservation include 0 or a value greater than 0 that is evenly divisible by 1024. The supported values for a limit include -1, 0, or a value greater than 0 that is evenly divisible by 1024.

### Troubleshooting Failed Shut Down VM Action Status

A shutdown VM action task has a `Failed` status in the Recent Task list.

The Shut Down VM action did not run successfully.

The Recent Tasks list indicates that a Shut Down VM action has a task status of `Failed`. When you evaluate the Messages list for the selected job, you see `Failure: Shut down confirmation timeout`.
The shutdown process involves shutting down the guest operating system and powering off the virtual machine. The wait time is 120 seconds to shut down the guest operating system. If the guest operating system does not shut down in this time, the action fails because the shutdown action is not confirmed.

1. To determine why the guest operating system did not shut down in the allotted time, check its status in vCenter.

#### Troubleshooting VMware Tools Not Running for a Shutdown VM Action Status

A Shutdown VM action task has a `Failed` status in the Recent Task list and the Message indicates that VMware Tools were required.

The Shutdown VM action did not run successfully.

The Recent Tasks list indicates that a Shutdown VM action has a tasks status of `Failed`. When you evaluate the Messages list for the selected job, you see `VMware Tools: Not running (Not installed).`
The Shutdown VM action requires that VMware Tools is installed and running on the target virtual machines. If you ran the action on more than one object, then VMware Tools was not installed, or installed but not running, on at least one of the virtual machines.

1. In the vCenter instance that manages the virtual machine that failed to run the action, install and start VMware Tools on the affected virtual machines.

#### Troubleshooting Failed Delete Unused Snapshots Action Status

A Delete Unused Snapshots action task has a `Failed` status in the Recent Task list.

The Delete Unused Snapshots action did not run successfully.

The Recent Tasks list indicates that a Delete Unused Snapshots action has a task status of `Failed`. When you evaluate the Messages list for the selected job, you see this message.

```
Remove snapshot failed, response wait expired after:[120] seconds,

     unable to confirm removal.
```

The delete snapshot process involves waiting for access to datastores. The wait time is 600 seconds to access the datastore and delete the snapshot. If the delete request is not passed to the datastore in that time, the action does not finish the delete snapshot action.

1. To determine if the snapshot was deleted, check its status in vCenter .
2. If it was not, submit the delete snapshot request at a different time.

## Metric, Property, and Alert Definitions

VMware Aria OperationsVMware Cloud Foundation Operations provides definitions for the metrics, properties, and alerts defined on objects in your environment.

# Metric Definitions in VMware Aria OperationsVMware Cloud Foundation Operations

Metric definitions provide an overview of how the metric value is calculated or derived. If you understand the metric, you can better tune VMware Aria OperationsVMware Cloud Foundation Operations to display results that help you to manage your environment.

VMware Aria OperationsVMware Cloud Foundation Operations collects data from objects in your environment. Each piece of data collected is called a metric observation or value. VMware Aria OperationsVMware Cloud Foundation Operations uses the VMware vCenter adapter to collect raw metrics. VMware Aria OperationsVMware Cloud Foundation Operations uses the VMware Aria OperationsVMware Cloud Foundation Operations adapter to collect self-monitoring metrics. In addition to the metrics it collects, VMware Aria OperationsVMware Cloud Foundation Operations calculates capacity metrics, badge metrics, and metrics to monitor the health of your system.

VMware Aria OperationsVMware Cloud Foundation Operations collects data from objects in your environment. Each piece of data collected is called a metric observation or value. VMware Aria OperationsVMware Cloud Foundation Operations uses the VMware vCenter® adapter to collect raw metrics. VMware Aria OperationsVMware Cloud Foundation Operations uses the VMware Aria OperationsVMware Cloud Foundation Operations adapter to collect self-monitoring metrics. In addition to the metrics it collects, VMware Aria OperationsVMware Cloud Foundation Operations calculates capacity metrics, badge metrics, and metrics to monitor the health of your system.

All metric definitions are provided. The metrics reported on your system depend on the objects in your environment. You can use metrics to help troubleshoot problems. See Troubleshooting with the Metrics Tab.

All metric definitions are provided. The metrics reported on your system depend on the objects in your environment. You can use metrics to help troubleshoot problems.

### Metrics for vCenter Components

VMware Aria OperationsVMware Cloud Foundation Operations connects to VMware vCenter® instances through the vCenter adapter to collect metrics for vCenter components and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

vCenter components are listed in the `describe.xml` file for the vCenter adapter. The following example shows sensor metrics for the host system in the `describe.xml` file.

```
<ResourceGroup instanced="false" key="Sensor" nameKey="1350" validation="">

    <ResourceGroup instanced="false" key="fan" nameKey="1351" validation="">

        <ResourceAttribute key="currentValue" nameKey="1360" dashboardOrder="1"
dataType="float"  defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal=""
minVal=""  unit="percent"/>

        <ResourceAttribute key="healthState" nameKey="1361" dashboardOrder="1"
dataType="float"  defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal=""
minVal="" />

    </ResourceGroup>

    <ResourceGroup instanced="false" key="temperature" nameKey="1352" validation="">

        <ResourceAttribute key="currentValue" nameKey="1362" dashboardOrder="1"
dataType="float"  defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal=""
minVal="" />

        <ResourceAttribute key="healthState" nameKey="1363" dashboardOrder="1"
dataType="float"  defaultMonitored="false"   isDiscrete="false" isRate="false" maxVal=""
minVal="" />

    </ResourceGroup>
```

```
</ResourceGroup>
```

Each `ResourceAttribute` element includes the name of a metric that appears in the UI and is documented as a Metric Key.

**Table 406: Sensor Metrics for Host System Cooling**

| Metric Key | Metric Name | Description |
|---|---|---|
| Sensor\|fan\|currentValue | Speed | Fan speed. |
| Sensor\|fan\|healthState | Health State | Fan health state. |
| Sensor\|temperature\|currentValue | Temperature | Host system temperature. |
| Sensor\|temperature\|healthState | Health State | Host system health state. |

### vSphere Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects CPU use, disk, memory, network, and summary metrics for objects in the vSphere world.

Capacity metrics can be calculated for vSphere world objects. See Capacity Analytics Generated Metrics.

#### vSphere World Super Metrics for ROI Dashboard

vSphere world super metrics provide information about the new metrics added to the ROI dashboard.

| Metric Name | Description |
|---|---|
| Cost\|Total Cost of Ownership | This metric shows the total cost cost of ownership with potential savings and optimizations.<br>Key: cost\|total_aggregated_cost |
| Online Capacity Analytics Capacity Remaining Profiles | This metric shows the VMs remaining based on the average VM profile.<br>Key: OnlineCapacityAnalytics\|capacityRemainingProfile |
| Cost\|Server Hardware(Owned) Cost | This metric shows the sum of server hardware depreciated cost having purchase type as Owned across all the vCenters.<br>Key: cost\|total_serverHardware_owned_cost |
| Cost\|Server Hardware(Leased) Cost | This metric shows the sum of server hardware depreciated cost having purchase type as Leased across all the vCenters.<br>Key: cost\|total_serverHardware_leased_cost |
| Cost\|Host OS License cost | This metric shows the sum of host OS license cost across all the vCenters.<br>Key: cost\|total_hostOsl_cost |
| Cost\|Network Cost | This metric shows the sum of network cost across all the vCenters.<br>Key: cost\|total_network_cost |
| Cost\|Maintenance Cost | This metric shows the sum of maintenance cost across all the vCenters.<br>Key: cost\|total_maintenance_cost |
| Cost\|Server Labor Cost | This metric shows the sum of server labor cost across all the vCenters.<br>Key: cost\|total_serverLabor_cost |
| Cost\|Facilities Cost | This metric shows the sum of facilities cost across all the vCenters.<br>Key: cost\|total_facilities_cost |
| Cost\|Additional Cost | This metric shows the sum of additional cost across all the vCenters.<br>Key: cost\|total_additional_cost |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Cost\|VM Direct Cost | This metric shows sum of direct Cost (VI labor + OS Labor) across all vCenters.<br>Key: cost\|total_vm_direct_cost |
| Cost\|Capacity Used Compute Cost | This metric displays the cost of the used compute capacity.<br>Key: cost\|capacity_used\|compute |
| Cost\|Capacity Remaining Compute Cost | This metric displays the cost of the remaining compute capacity.<br>Key: cost\|capacity_remaining\|compute |
| Cost\|Capacity Used Storage Cost | This metric displays the cost of the used storage capacity.<br>Key: cost\|capacity_used\|storage |
| Cost\|Capacity Remaining Storage Cost | This metric displays the cost of the remaining storage capacity.<br>Key: cost\|capacity_remaining\|storage |
| Cost\|Potential Savings Idle VMs | This metric displays the potential savings from Idle VMs.<br>Key: cost\|potential_savings\|idle_vms |
| Cost\|Potential Savings Powered Off VMs | This metric displays the potential savings from powered off VMs.<br>Key: cost\|potential_savings\|poweredOff_vms |
| Cost\|Potential Savings VM Snapshots | This metric displays the potential savings from VM snapshots.<br>Key: cost\|potential_savings\|vm_snapshots |
| Cost\|Potential Savings Orphaned Disks | This metric displays the potential savings from orphaned disks.<br>Key: cost\|potential_savings\|orphaned_disks |
| Cost\|Potential Savings Oversized VMs | This metric displays the potential savings from oversized VMs.<br>Key: cost\|potential_savings\|oversized_vms |
| Cost\|Potential Savings Cost Optimization Opportunities | This metric displays the potential savings from cost optimization opportunities.<br>Key: cost\|potential_savings\|cost_optimization_opportunities |
| Cost\|Total Cost of Ownership | This metric shows the total cost cost of ownership with potential savings and optimizations.<br>Key: cost\|potential_savings\|total_cost_of_ownership |
| Server Purchase Cost | This metric shows the server purchase cost.<br>Key: cost\|server_purchase_cost |
| Accumulated Depreciation | This metric displays the sum of the accumulated depreciation (Depreciation is calculated from the purchase date till current date) of servers across all the vCenters.<br>Key: cost\|accumulatedDepreciation |
| Remaining Depreciation | This metric displays the sum of the remaining depreciation (Remaining Depreciation is calculated from the current date till Depreicated year) of servers across all the vCenters.<br>Key: cost\|accumulatedDepreciation |
| Number of Fully Depreciated Servers | This metric displays the number of fully depreciated servers across all the vCenters.<br>Key: cost\|hardwareTotalCost |
| Reclaimed vCPUs from Idle VMs | This metric displays the number of reclaimable vCPUs from idle VMs.<br>Key: reclaimable\|idle_vms\|cpu |
| Reclaimed Memory from Idle VMs | This metric displays the amount of reclaimable memory from the idle VMs.<br>Key: reclaimable\|idle_vms\|mem |
| Reclaimed Disk Space from Idle VMs | This metric displays the amount of reclaimable disk space from the idle VMs.<br>Key: reclaimable\|idle_vms\|diskspace |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Reclaimed Disk Space from Powered Off VMs | This metric displays the amount of reclaimable disk space from the powered off VMs.<br>Key: reclaimable\|poweredOff_vms\|diskspace |
| Reclaimed Disk Space from VM Snapshots | This metric displays the amount of reclaimable disk space from the VM Snapshots.<br>Key: reclaimable\|vm_snapshots\|diskspace |
| Reclaimed Disk Space from Orphaned Disks | This metric displays the amount of reclaimable disk space from the orphaned disks.<br>Key: reclaimable\|orphaned_disk\|diskspace |
| Rightsize - vCPUs to Remove from Oversized VMs | This metric displays the number of vCPUs to remove from the oversized VMs.<br>Key: summary\|oversized\|vcpus |
| Rightsize - Memory to Remove from Oversized VMs | This metric displays the amount of memory to be removed from the oversized VMs.<br>Key: summary\|oversized\|memory |
| Rightsize - vCPUs to Add from Undersized VMs | This metric displays the number of vCPUs to be added from the undersized VMs.<br>Key: summary\|undersized\|vcpus |
| Rightsize - Memory to Add from Undersized VMs | This metric displays the amount of memory to be added from the undersized VMs.<br>Key: summary\|undersized\|memory |
| Total Storage Cost | This metric displays the sum of storage cost across all vCenters.<br>Key: cost\|totalCost |
| Total Potential Savings | This metric displays the sum of all the potential savings (Idle VMs + Powered off Vms + Snapshot + Orphaned Disks + Oversized VMs).<br>Key: reclaimable\|cost |
| **New vSphere Metrics Added for ROI Dashboard** | |
| Potential Savings from Oversized VMs | This metric displays the sum of all the potentials savings gained from oversized VMs across vcenters.<br>Key: cost\|reclaimableCost |
| Reclaimable Host Cost | This metric displays the reclaimable host cost based on the recommended size.<br>Key: cost\|potential_savings\|total_reclaimable_host_cost |
| Cost\|Potential Increase\|Undersized VMs Cost | This metric displays the rightsizing value for the undersized VMs.<br>Key: cost\|potential_increase\|undersized_vms |
| Cost\|Realized Savings\|Total Realized Savings | This metric displays the total realize savings for VMs across all vCenters.<br>Key: cost\|realized_savings\|total_realized_savings |
| Cost\|Realized Savings\|Idle Savings | This metric displays the total realized savings for idle VMs across all vCenters.<br>Key: cost\|realized_savings\|realized_idle_savings |
| Cost\|Realized Savings\|Powered Off Savings | This metric displays the total realized savings for powered off VMs across all vCenters.<br>Key: cost\|realized_savings\|realized_poweredOff_savings |
| Cost\|Realized Savings\|Snapshot Space Savings | This metric displays the total realized savings for snapshot space across all vCenters.<br>Key: cost\|realized_savings\|realized_snapshotSpace_savings |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Cost\|Realized Savings\|Oversized Savings | This metric displays the oversized savings across all vCenters.<br>Key: cost\|realized_savings\|realized_oversized_savings |
| Cost\|Realized Savings\|Orphaned Disk Space Savings | This metric displays the amount of disk space saved by orphaned disks across all vCenters.<br>Key: cost\|realized_savings\|realized_orphanedDiskSpace_savings |
| Cost\|Realized Savings\|Reclaimable Host Savings | This metric displays the amount of reclaimable host savings across all vCenters.<br>Key: cost\|realized_savings\|realized_reclaimableHost_savings |
| Compute Realized\|vCPUs from Oversized VMs | This metric displays the number of vCPUs realized across all vCenters.<br>Key: compute_realized\|realized_oversized_vcpus |
| Compute Realized\|Memory from Oversized VMs | This metric displays the amount of memory realized from oversized VMs across all vCenters.<br>Key: compute_realized\|realized_oversized_mem |
| Realized Potential Memory Consumed from Oversized VMs | This metric displays the potential memory consumed from oversized VMs across all vCenters.<br>Key: realized\|realizedPotentialMemConsumed |
| Total Number Of Reclaimable Hosts | This metric displays the total number of reclaimable hosts across all vCenters.<br>Key: metric=cost\|reclaimableHostCost |
| Compute Realized\|vCPUs from Idle VMs | This metric displays the realized vCPUs from idle VMs across all vCenters.<br>Key: compute_realized\|realized_idle_vcpus |
| Compute Realized\|Memory from Idle VMs | This metric displays the amount of memory realized from idle VMs across all vCenters.<br>Key: compute_realized\|realized_idle_mem |
| Disk Space Realized\|Idle VMs | This metric displays the amount of disk space realized from idle VMs across all vCenters.<br>Key: storage_realized\|realized_idle_diskSpace |
| Disk Space Realized\|Powered Off VMs | This metric displays the amount of disk space realized from powered off VMs across all vCenters.<br>Key: storage_realized\|realized_poweredOff_diskSpace |
| Disk Space Realized\|VM Snapshots | This metric displays the amount of disk space realized from VM snapshots across all vCenters.<br>Key: storage_realized\|realized_snapshotSpace |
| Disk Space Realized\|Orphaned Disks | This metric displays the amount of disk space realized from orphaned disks across all vCenters.<br>Key: storage_realized\|realized_orphaned_diskSpace |

**CPU Usage Metrics**

CPU usage metrics provide information about CPU use.

| Metric Name | Description |
|---|---|
| CPU\|Capacity usage | CPU usages as a percent during the interval.<br>Key: cpu\|capacity_usagepct_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| CPU\|CPU contention(%) | This metric shows the percentage of time the VMs in the ESXi hosts are unable to run because they are contending for access to the physical CPUs. The number shown is the average number for all VMs. This number is lower than the highest number experienced by the VM most impacted by CPU contention.<br><br>Use this metric to verify if the host can serve all its VMs efficiently. Low contention means that the VM can access everything it demands to run smoothly. It means that the infrastructure is providing good service to the application team.<br><br>When using this metric, ensure that the number is within your expectation. Look at both the relative number and the absolute number. Relative means a drastic change in value, meaning that the ESXi is unable to serve the VMs. Absolute means that the real value itself is high. Investigate why the number is high. One factor that impacts this metric is CPU Power Management. If CPU Power Management clocks down the CPU speed from 3 GHz to 2 GHz, the reduction in speed is accounted for because it shows that the VM is not running at full speed.<br><br>This metric is calculated in the following way: cpu\|capacity_contention / (200 * summary\|number_running_vcpus)<br><br>Key: cpu\|capacity_contentionPct |
| CPU\|Demand (%) | This metric shows the amount of CPU resources a virtual machine might use if there were no CPU contention or CPU limit. This metric represents the average active CPU load for the past five minutes.<br><br>Keep this number below 100% if you set the power management to maximum.<br><br>This metric is calculated in the following way: ( cpu.demandmhz / cpu.capacity_provisioned)*100<br><br>Key: cpu\|demandPct |
| CPU\|Demand (MHz) | This metric shows the amount of CPU resources a virtual machine might use if there were no CPU contention or CPU limit.<br>Key: cpu\|demandmhz |
| CPU\|Demand | CPU demand in megahertz.<br>Key: cpu\|demand_average |
| CPU\|IO wait | IO wait (ms).<br>Key: cpu\|iowait |
| CPU\|number of CPU Sockets | Number of CPU sockets.<br>Key: cpu\|numpackages |
| CPU\|Overall CPU Contention | Overall CPU contention in milliseconds.<br>Key: cpu\|capacity_contention |
| CPU\|Provisioned Capacity (MHz) | Capacity in MHz of the physical CPU cores.<br>Key: cpu\|capacity_provisioned |
| CPU\|Provisioned vCPU(s) | Number of provisioned CPU cores. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: cpu\|corecount_provisioned |
| CPU\|Reserved Capacity (MHz) | Total CPU capacity reserved by virtual machines.<br>Key: cpu\|reservedCapacity_average |
| CPU\|Usage (MHz) | CPU usages, as measured in megahertz, during the interval.<br>• VM - Amount of actively used virtual CPU. This is the host's view of the CPU usage, not the guest operating system view.<br>• Host - Sum of the actively used CPU of all powered on virtual machines on a host. The maximum possible value is the frequency of the two processors multiplied by the number of processors. For example, if you have a host with four 2 GHz CPUs running a virtual machine that is using 4000 MHz, the host is using two CPUs completely: 400 / (4 2000) = 0.50<br>Key: cpu\|usagemhz_average |
| CPU\|Wait | Total CPU time spent in wait state. The wait total includes time spent in the CPU Idle, CPU Swap Wait, and CPU I/O Wait states.<br>Key: cpu\|wait |
| CPU\|Workload (%) | Percent of workload<br>Key: cpu\|workload |

**Memory Metrics**

Memory metrics provide information about memory use and allocation.

| Metric Name | Description |
|---|---|
| mem\|Contention (%) | This metric shows the percentage of time VMs are waiting to access swapped memory.<br><br>Use this metric to monitor ESXi memory swapping. A high value indicates that the ESXi is running low on memory, and a large amount of memory is being swapped.<br><br>Key: mem\|host_contentionPct |
| mem\|Machine Demand (KB) | Host memory demand in kilobytes.<br>Key: mem\|host_demand |
| mem\|Provisioned Memory | Provisioned host memory in kilobytes.<br>Key: mem\|host_provisioned |
| mem\|Reserved Capacity (KB) | Total amount of memory reservation used by powered-on virtual machines and vSphere services on the host.<br>Key: mem\|reservedCapacity_average |
| mem\|Usable Memory (KB) | Usable host memory in kilobytes.<br>Key: mem\|host_usable |
| mem\|Host Usage (KB) | Host memory use in kilobytes.<br>Key: mem\|host_usage |
| mem\|Usage/Usable (%) | Memory usage as percentage of total configured or available memory.<br>Key: mem\|host_usagePct |
| mem\|Workload (%) | Percent of workload.<br>Key: mem\|workload |

**Network Metrics**

Network metrics provide information about network performance.

| Metric Name | Description |
|---|---|
| net\|Packets Dropped (%) | This metric shows the percentage of received and transmitted packets dropped in the collection interval.<br>Use this metric to monitor the reliability and performance of the ESXi network. A high value indicates that the network is not reliable and performance decreases.<br><br>Key: net\|droppedPct |
| net\|Usage Rate (KB per second) | Sum of the data transmitted and received for all of the NIC instances of the host or virtual machine.<br>Key: net\|usage_average |
| net\|Workload (%) | Percent of workload.<br>Key: net\|workload |

**Disk Metrics**

Disk metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| disk\|Total IOPS | Average number of commands issued per second during the collection cycle.<br>Key: disk\|commandsAveraged_average |
| disk\|Usage Rate (KB per second) | Average of the sum of the data read and written for all of the disk instances of the host or virtual machine.<br>Key: disk\|usage_average |
| disk\|Workload (%) | Percent of workload.<br>Key: disk\|workload |

**Summary Metrics**

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| summary\|Number of Running Hosts | Number of running hosts.<br>Key: summary\|number_running_hosts |
| summary\|Number of Running VMs | This metric shows the number of running VMs at a given point in time. The data is sampled every five minutes.<br><br>A large number of running VMs might be a reason for CPU or memory spikes because more resources are used in the host. The number of running VMs gives you a good indicator of how many requests the ESXi host must juggle. Powered off VMs are not included because they do not impact ESXi performance. A change in the number of running VMs can contribute to performance problems. A high number of running VMs in a |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | host also means a higher concentration risk, because all the VMs fail if the ESXi crashes.<br><br>Use this metric to look for a correlation between spikes in the running VMs and spikes in other metrics such as CPU contention, or memory contention.<br><br>Key: summary\|number_running_vms |
| summary\|Number of Clusters | Total number of clusters.<br>Key: summary\|total_number_clusters |
| summary\|Total Number of Datastores | Total number of datastores.<br>Key: summary\|total_number_datastores |
| summary\|Number of Hosts | Total number of hosts.<br>Key: summary\|total_number_hosts |
| summary\|Number of VMs | Total number of virtual machines.<br>Key: summary\|total_number_vms |
| summary\|Total Number of Datacenters | Total number of data centers.<br>Key: summary\|total_number_datacenters |
| summary\|Number VCPUs on Powered on VMs | Number of virtual CPUs on powered-on virtual machines.<br>Key: summary\|number_running_vcpus |
| summary\|Average Running VM Count per Running Host | Average running virtual machine count per running host.<br>Key: summary\|avg_vm_density |
| summary\|Number of Reclaimable Hosts | Displays the number of reclaimable hosts.<br>Key: summary\|total_number_reclaimable_hosts |

**Virtual Machine Operations Metrics for vSphere World**

VM operations metrics provide information about the actions performed on VMs. The following are some important points you must know about VM operation metrics for vSphere World.

- VM operations metrics is not collected for custom data centers.
- If you edit a VM settings and do not perform any action, still it is considered as VM reconfigure operation.
- During Revert Snapshot, VMs are powered-off, but this operation is not counted under VM Power-off metric.
- Adding ESXi with VMs is not counted under VM Create metric.
- Removing ESXi with VMs is not coutned under VM Remove metric.
- VM hardstop operation is not counted under VM Power Off metric.

| Metric Name | Description |
|---|---|
| Inventory | |
| VM Clone | This metric displays the number of clone operations on the virtual machine.<br>Key: Inventory\|VM Clone |
| VM Create | This metric displays the number of create operations on the virtual machine.<br>Key: Inventory\|VM Create |
| VM Delete | This metric displays the number of delete operations on the virtual machine. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: Inventory\|VM Delete |
| VM Reconfigure | This metric displays the number of reconfigure operations on the virtual machine.<br>Key: Inventory\|VM Reconfigure |
| VM Register | This metric displays the number of register operations on the virtual machine.<br>Key: Inventory\|VM Register |
| VM Template Deploy | This metric displays the number templates deployed on the virtual machine.<br>Key: Inventory\|VM Template Deploy |
| VM Unregister | This metric displays the number of unregister operations on the virtual machine.<br>Key: Inventory\|VM Unregister |
| Location | |
| Storage vMotion | This metric displays the number of migrations with vMotion (datastore change operations for Powered-on VMs).<br>Key: Location\|Storage vMotion |
| VM Datastore Change (powered-off VMs) | This metric displays the number of datastore change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-off VMs) | This metric displays the number of host and datastore change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Host and Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-on VMs) | This metric displays the number of host and datastore change operations, for powered-on and suspended virtual machines.<br>Key: Location\|VM Host and Datastore Change (powered-on VMs) |
| VM Host Change (powered-off VMs) | This metric displays the number of host change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Host Change (powered-off VMs) |
| vMotion | This metric displays the number of migrations with vMotion (host change operations for powered-on VMs).<br>Key: Location\|vMotion |
| State | |
| VM Guest Reboot | This metric displays the number of reboot operations on the virtual machine guest.<br>Key: State\|VM Guest Reboot |
| VM Guest Shutdown | This metric displays the number of shutdown operations on the virtual machine guest.<br>Key: State\|VM Guest Shutdown |
| VM Power Off | This metric displays the number of power-off operations on the virtual machine.<br>Key: State\|VM Power Off |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| VM Power On | This metric displays the number of power-on operations on the virtual machine.<br>Key: State\|VM Power On |
| VM Reset | This metric displays the number of reset operations on the virtual machine guest.<br>Key: State\|VM Reset |
| VM Standby Guest | This metric displays the number of standby operations on the virtual machine guest.<br>Key: State\|VM Standby Guest |
| VM Suspend | This metric displays the number of suspend operations on the virtual machine.<br>Key: State\|VM Suspend |

### vCenter Server Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects CPU use, disk, memory, network, and summary metrics for vCenter Server system objects.

vCenter Server metrics include capacity and badge metrics. See definitions in:
- Capacity Analytics Generated Metrics
- Badge Metrics

### CPU Usage Metrics

CPU usage metrics provide information about CPU use.

| Metric Name | Description |
|---|---|
| Capacity Usage (%) | Percent capacity used.<br>Key: cpu\|capacity_usagepct_average |
| CPU Contention (%) | Percent CPU contention.<br>Key: cpu\|capacity_contentionPct |
| Demand (%) | Percent demand.<br>Key: cpu\|demandPct |
| Demand (MHz) | CPU utilization level based on descendant virtual machines utilization. This Includes reservations, limits, and overhead to run the virtual machines.<br>Key: cpu\|demandmhz |
| Demand | CPU Demand.<br>Key: cpu\|demand_average |
| IO Wait (ms) | IO wait time in milliseconds.<br>Key: cpu\|iowait |
| Number of CPU Sockets | Number of CPU sockets.<br>Key: cpu\|numpackages |
| Overall CPU Contention (ms) | Overall CPU contention in milliseconds.<br>Key: cpu\|capacity_contention |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Provisioned Capacity (MHz) | Provisioned capacity in megahertz.<br>Key: cpu\|capacity_provisioned |
| Provisioned vCPU | Number of provisioned virtual CPU cores.<br>Key: cpu\|corecount_provisioned |
| Reserved Capacity (MHz) | Sum of the reservation properties of the immediate children of the host's root resource pool.<br>Key: cpu\|reservedCapacity_average |
| Usage (MHz) | Average CPU use in megahertz.<br>Key: cpu\|usagemhz_average |
| Wait (ms) | CPU time spent on the idle state.<br>Key: cpu\|wait |
| Overhead | Amount of CPU that is overhead.<br>Key: cpu\|overhead_average |
| Demand without overhead | Value of demand excluding any overhead.<br>Key: cpu\|demand_without_overhead |
| Provisioned Capacity | Provisioned capacity (MHz).<br>Key: cpu\|vm_capacity_provisioned |
| Total Capacity (MHz) | Total CPU resources configured on the descendant ESXi hosts.<br>Key: cpu\|capacity_provisioned |
| Usable Capacity (MHz) | The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services.<br>Key: cpu\|haTotalCapacity_average |

**Datastore Metrics**

Datastore metrics provide information about the datastore.

| Metric Name | Description |
|---|---|
| Outstanding IO requests | OIO for datastore.<br>Key: datastore\|demand_oio |
| Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: datastore\|numberReadAveraged_average |
| Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: datastore\|numberWriteAveraged_average |
| Read Throughput (KBps) | Amount of data read in the performance interval.<br>Key: datastore\|read_average |
| Write Throughput (KBps) | Amount of data written to disk in the performance interval.<br>Key: datastore\|write_average |

**Disk Metrics**

Disk metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| Total IOPS | Average number of commands issued per second during the collection cycle.<br>Key: disk\|commandsAveraged_average |
| Total Latency (ms) | Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Device Command Latency and Physical Device Command Latency metrics.<br>Key: disk\|totalLatency_average |
| Total Throughput (KBps) | Average of the sum of the data read and written for all the disk instances of the host or virtual machine.<br>Key: disk\|usage_average |
| Total queued outstanding operations | Sum of queued operations and outstanding operations.<br>Key: disk\|sum_queued_oio |
| Max Observed OIO | Max observed IO for a disk.<br>Key: disk\|max_observed |

**Disk Space Metrics**

Disk space metrics provide information about disk space use.

| Metric Name | Description |
|---|---|
| Total disk space used (KB) | Total disk space used on all datastores visible to this object.<br>Key: diskspace\|total_usage |
| Total disk space (KB) | Total disk space on all datastores visible to this object.<br>Key: diskspace\|total_capacity |
| Total provisioned disk space (KB) | Total provisioned disk space on all datastores visible to this object.<br>Key: diskspace\|total_provisioned |
| Utilization (GB) | Storage space used on the connected vSphere Datastores.<br>Key: diskspace\|total_usage |
| Total Capacity (GB) | Total storage space available on the connected vSphere datastores.<br>Key: diskspace\|total_capacity |

**Memory Metrics**

Memory metrics provide information about memory use and allocation.

| Metric Name | Description |
|---|---|
| Contention (%) | Percent host memory contention.<br>Key: mem\|host_contentionPct |
| Machine Demand (KB) | Host memory demand in kilobytes.<br>Key: mem\|host_demand |
| ESX System Usage | Memory usage by the VMkernel and ESX user-level services.<br>Key: mem\|host_systemUsage |
| Provisioned Memory (KB) | Provisioned host memory in kilobytes.<br>Key: mem\|host_provisioned |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Reserved Capacity (KB) | Sum of the reservation properties of the immediate children of the host's root resource pool.<br>Key: mem\|reservedCapacity_average |
| Usable Memory (KB) | Usable host memory in kilobytes.<br>Key: mem\|host_usable |
| Host Usage (KB) | Host memory use in kilobytes.<br>Key: mem\|host_usage |
| Usage/Usable (%) | Percent host memory used.<br>Key: mem\|host_usagePct |
| Contention (KB) | Host contention in kilobytes.<br>Key: mem\|host_contention |
| VM Overhead (KB) | Memory overhead reported by host.<br>Key: mem\|overhead_average |
| Utilization (KB) | Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines.<br>Key: mem\|total_need |
| Total Capacity (KB) | Total physical memory configured on descendant ESXi hosts.<br>Key: mem\|host_provisioned |
| Usable Capacity (KB) | The usable memory resources available for the virtual machines after considering reservations for vSphere HA and other vSphere services.<br>Key: mem\|haTotalCapacity_average |

**Network Metrics**

Network metrics provide information about network performance.

| Metric Name | Description |
|---|---|
| Packets Dropped (%) | Percent network packets dropped.<br>Key: net\|droppedPct |
| Total Throughput (KBps) | Sum of the data transmitted and received for all of the NIC instances of the host or virtual machine.<br>Key: net\|usage_average |
| Packets Received | Number of packets received in the performance interval.<br>Key: net\|packetsRx_summation |
| Packets Transmitted | Number of packets transmitted in the performance interval.<br>Key: net\|packetsTx_summation |
| Received Packets Dropped | Number of received packets dropped in the performance interval.<br>Key: net\|droppedRx_summation |
| Transmitted Packets Dropped | Number of transmitted packets dropped in the performance interval.<br>Key: net\|droppedTx_summation |
| Data Transmit Rate (KBps) | Average amount of data transmitted per second.<br>Key: net\|transmitted_average |
| Data Receive Rate (KBps) | Average amount of data received per second.<br>Key: net\|received_average |

### Summary Metrics

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| Number of Running Hosts | Number of hosts that are on.<br>Key: summary\|number_running_hosts |
| Number of Running VMs | Number of virtual machines that are on.<br>Key: summary\|number_running_vms |
| Number of Clusters | Total number of clusters.<br>Key: summary\|total_number_clusters |
| Total Number of Datastores | Total number of datastores.<br>Key: summary\|total_number_datastores |
| Number of Hosts | Total number of hosts.<br>Key: summary\|total_number_hosts |
| Number of VMs | Total number of virtual machines.<br>Key: summary\|total_number_vms |
| Maximum Number of VMs | Maximum number of virtual machines.<br>Key: summary\|max_number_vms |
| Workload Indicator (%) | Percent workload indicator.<br>Key: summary\|workload_indicator |
| Total Number of data centers | Total number of data centers.<br>Key: summary\|total_number_datacenters |
| Number of Cores on Powered On Hosts | Number of cores on powered-on hosts.<br>Key: summary\|number_powered_on_cores |
| Number VCPUs on Powered on VMs | Number of virtual CPUs on powered-on virtual machines.<br>Key: summary\|number_running_vcpus |
| Average Running VM Count per Running Host | Average running virtual machine count per running host.<br>Key: summary\|avg_vm_density |
| VC Query Time (ms) | vCenter Server query time in milliseconds.<br>Key: summary\|vc_query_time |
| Derived Metrics Computation Time (ms) | Derived metrics computation time in milliseconds.<br>Key: summary\|derived_metrics_comp_time |
| Number of objects | Number of objects.<br>Key: summary\|number_objs |
| Number of VC Events | Number of vCenter Server events.<br>Key: summary\|number_vc_events |
| Number of SMS Metrics | Number of SMS metrics.<br>Key: summary\|number_sms_metrics |
| Collector Memory Usage (MB) | Collector memory use in megabytes.<br>Key: summary\|collector_mem_usage |

### Virtual Machine Operations Metrics for vCenter Server

VM operations metrics provide information about the actions performed on VMs. The following are some important points you must know about VM operation metrics for vCenter Server.

- VM operations metrics is not collected for custom data centers.
- If you edit a VM settings and do not perform any action, still it is considered as VM reconfigure operation.
- During Revert Snapshot, VMs are powered-off, but this operation is not counted under VM Power-off metric.

- Adding ESXi with VMs is not counted under VM Create metric.
- Removing ESXi with VMs is not coutned under VM Remove metric.
- VM hardstop operation is not counted under VM Power Off metric.

| Metric Name | Description |
|---|---|
| Inventory | |
| VM Clone | This metric displays the number of clone operations on the virtual machine.<br>Key: Inventory\|VM Clone |
| VM Create | This metric displays the number of create operations on the virtual machine.<br>Key: Inventory\|VM Create |
| VM Delete | This metric displays the number of delete operations on the virtual machine.<br>Key: Inventory\|VM Delete |
| VM Reconfigure | This metric displays the number of reconfigure operations on the virtual machine.<br>Key: Inventory\|VM Reconfigure |
| VM Register | This metric displays the number of register operations on the virtual machine.<br>Key: Inventory\|VM Register |
| VM Template Deploy | This metric displays the number templates deployed on the virtual machine.<br>Key: Inventory\|VM Template Deploy |
| VM Unregister | This metric displays the number of unregister operations on the virtual machine.<br>Key: Inventory\|VM Unregister |
| Location | |
| Storage vMotion | This metric displays the number of migrations with vMotion (datastore change operations for Powered-on VMs).<br>Key: Location\|Storage vMotion |
| VM Datastore Change (powered-off VMs) | This metric displays the number of datastore change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-off VMs) | This metric displays the number of host and datastore change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Host and Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-on VMs) | This metric displays the number of host and datastore change operations, for powered-on and suspended virtual machines.<br>Key: Location\|VM Host and Datastore Change (powered-on VMs) |
| VM Host Change (powered-off VMs) | This metric displays the number of host change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Host Change (powered-off VMs) |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| vMotion | This metric displays the number of migrations with vMotion (host change operations for powered-on VMs).<br>Key: Location\|vMotion |
| State | |
| VM Guest Reboot | This metric displays the number of reboot operations on the virtual machine guest.<br>Key: State\|VM Guest Reboot |
| VM Guest Shutdown | This metric displays the number of shutdown operations on the virtual machine guest.<br>Key: State\|VM Guest Shutdown |
| VM Power Off | This metric displays the number of power-off operations on the virtual machine.<br>Key: State\|VM Power Off |
| VM Power On | This metric displays the number of power-on operations on the virtual machine.<br>Key: State\|VM Power On |
| VM Reset | This metric displays the number of reset operations on the virtual machine guest.<br>Key: State\|VM Reset |
| VM Standby Guest | This metric displays the number of standby operations on the virtual machine guest.<br>Key: State\|VM Standby Guest |
| VM Suspend | This metric displays the number of suspend operations on the virtual machine.<br>Key: State\|VM Suspend |

**Disabled Metrics**

The following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

You can enable these metrics in the Policy workspace. For more information, see Metrics and Properties Details.

| Metric Name | Description |
|---|---|
| Max Observed Number of Outstanding IO Operations | Maximum observed number of outstanding IO operations.<br>Key: datastore\|maxObserved_OIO |
| Max Observed Read Rate | Max observed rate of reading data from the datastore.<br>Key: datastore\|maxObserved_Read |
| Max Observed Reads per second | Max observed average number of read commands issued per second during the collection interval.<br>Key: datastore\|maxObserved_NumberRead |
| Max Observed Writes per second | Max observed average number of write commands issued per second during the collection interval. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: datastore\|maxObserved_NumberWrite |
| Max Observed Write Rate | Max observed rate of writing data from the datastore.<br>Key: datastore\|maxObserved_Write |
| Max Observed Throughput (KBps) | Max observed rate of network throughput.<br>Key: net\|maxObserved_KBps |
| Max Observed Transmitted Throughput (KBps) | Max observed transmitted rate of network throughput.<br>Key: net\|maxObserved_Tx_KBps |
| Max Observed Received Throughput (KBps) | Max observed received rate of network throughput.<br>Key: net\|maxObserved_Rx_KBps |

**Virtual Machine Metrics**

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, CPU use, memory, datastore, disk, virtual disk, guest file system, network, power, disk space, storage, and summary metrics for virtual machine objects.

**Guest Operating System Metrics**

Guest Operating System metrics provide information about the new metrics added to the Guest Operating System

| Metric Name | Description |
|---|---|
| Guest\|Peak Guest OS Page-out/rate within collection cycle | This metric displays the highest memory page-out rate reported by guest operating system, measured as peak of any 20-second average during the collection interval.<br>**NOTE**<br>To collect the guest metrics, ensure that VM Tools is installed and up and running on virtual machine, on vCenter server.<br>Key: guest\|20_sec_peak_page.outRate_latest |

**Metrics for ROI Dashboard**

Virtual machine metrics provide information about the new metrics added to the ROI dashboard.

| Metric Name | Description |
|---|---|
| Potential Memory Consumed Reclaimable(GB) | This metric displays the sum of all the reclaimble consumed memory for the virtual machine. |
| Potential CPU Usage Increase (GHz) | This metric displays the potential increase in CPU usage for the virtual machine. |
| Potential Memory Usage Increase (GB) | This metric displays the potential increase in memory usage for the virtual machine. |
| Potential Savings | This metric displays the sum of all the potential savings (Idle VMs + Powered off Vms + Snapshot + Orphaned Disks + Oversized VMs). |
| Potential Cost Increase | This metric displays the potential increase in costs associated with the virtual machine. |

## Application Level Cost Roll up Metrics

The application-level cost roll up in VMware Aria OperationsVMware Cloud Foundation Operations includes few additional metrics at the application level. VMware Aria OperationsVMware Cloud Foundation Operations has introduced Business Application as a new object, the business application object can have Tiers and Applications as its children. The cost roll up option allows you to aggregate all VM costs associated with the application and tiers and publish them at VM level.

**Table 407: Application Level Cost Roll up Metrics**

| Metric Name | Description |
|---|---|
| Monthly Effected Projected Total Cost | This metric displays the effective virtual machine cost projected for the full month. |
| Effective MTD Cost | This metric displays the month to date effective application cost for the selected virtual machine. |
| Effective Daily Cost | This metric displays the effective daily cost of the application associated with the virtual machine. |

## Configuration Metrics for Virtual Machines

Configuration metrics provide information about virtual machine configuration.

| Metric Name | Description |
|---|---|
| Config\|Thin Provisioned Disk | Thin Provisioned Disk.<br>Key: config\|hardware\|thin_Enabled |
| Config\|Number of CPUs | Number of CPUs for a Virtual Machine.<br>From VMware Cloud Foundation Operations 6.7 and onwards, this metric is measured in vCPUs instead of cores.<br><br>Key: config\|hardware\|num_Cpu |
| Config\|Disk Space | Disk space metrics.<br>Key: config\|hardware\|disk_Space |

## CPU Usage Metrics for Virtual Machines

CPU usage metrics provide the information about CPU use.

| Metric Name | Description |
|---|---|
| CPU\|Other Wait (ms) | CPU time spent waiting for IO.<br>Key: cpu\|otherwait |
| CPU\|Overall CPU Contention (ms) | The amount of time the CPU cannot run due to contention.<br>Key: cpu\|capacity_contention |
| CPU\|Reservation Used | CPU Reservation Used.<br>Key: cpu\|reservation_used |
| CPU\|Effective Limit | CPU Effective Limit.<br>Key: cpu\|effective_limit |
| CPU\|Other Wait (%) | Percentage Other Wait.<br>Key: cpu\|otherwaitPct |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| CPU\|Swap wait (%) | Percentage swap waits for CPU.<br>Key: cpu\|swapwaitPct |
| CPU\|Wait (%) | Percentage of the total CPU time spent in wait state.<br>Key: cpu\|waitPct |
| CPU\|System (%) | Percentage CPU time spent on system processes.<br>Key: cpu\|systemSummationPct |
| CPU\|Capacity entitlement (MHz) | CPU entitlement for the VM after considering all limits.<br>Key: cpu\|capacity_entitlement |
| CPU\|Capacity Demand Entitlement (% | Percent capacity demand entitlement.<br>Key: cpu\|capacity_demandEntitlementPct |
| CPU\|CPU Contention (%) | CPU contention as a percentage of 20-second collection interval.<br>Key: cpu\|capacity_contentionPct |
| CPU\|Total Capacity | Provisioned CPU capacity in megahertz.<br>Key: cpu\|vm_capacity_provisioned |
| CPU\|Demand (MHz) | Total CPU resources required by the workloads on the virtual machine.<br>Key: cpu\|demandmhz |
| CPU\|Host demand for aggregation | Host demand for aggregation.<br>Key: cpu\|host_demand_for_aggregation |
| CPU\|Demand (ms) | The total CPU time that the VM might use if there was no contention.<br>Key: cpu\|demand_average |
| CPU\|Demand (%) | CPU demand as a percentage of the provisioned capacity.<br>Key: cpu\|demandPct |
| CPU\|Usage (%) | This metric indicates the percentage of CPU that was used out of all the CPU that was allocated to the VM. CPU usage can indicate when the VM is undersized.<br>Key: cpu\|usage_average |
| CPU\|Usage (MHz) | CPU use in megahertz.<br>Key: cpu\|usagemhz_average |
| CPU Workload % | This metric indicates the CPU workload % for the VM, the maximum threshold for this is 80% and the minimum threshold is 20%. If your Maximum line is constantly ~100% flat, you may have a runaway process. If this chart is below or less than 20% all the time for the entire month, then all the large VMs are oversized. This number must hover around 40%, indicating the sizing done was accurate. |
| CPU\|System (ms) | CPU time spent on system processes.<br>Key: cpu\|system_summation |
| CPU\|Ready (%) | This metric indicates the percentage of time in which the VM was waiting in line to use the CPU on the host. A large ready time for a VM indicates that the VM needed CPU resources but the infrastructure was busy |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | serving other VMs. A large ready time might indicate that the host is trying to serve too many VMs.<br><br>Whenever the CPU ready is larger than 10%, you should check if the host is overloaded, or if the VM really needs all the resources that were allocated to it.<br><br>Key: cpu\|readyPct |
| CPU\|Extra (ms) | Extra CPU time in milliseconds.<br>Key: cpu\|extra_summation |
| CPU\|Guaranteed (ms) | CPU time that is guaranteed for the virtual machine.<br>Key: cpu\|guaranteed_latest |
| CPU\|Co-stop (%) | Percentage of time the VM is ready to run, but is unable to due to co-scheduling constraints.<br>Key: cpu\|costopPct |
| CPU\|Latency | Percentage of time the VM is unable to run because it is contending for access to the physical CPUs.<br>Key: cpu\|latency_average |
| CPU\|Max Limited | Time the VM is ready to run, but is not run due to maxing out its CPU limit setting.<br>Key: cpu\|maxlimited_summation |
| CPU\|Overlap | Time the VM was interrupted to perform system services on behalf of that VM or other VMs.<br>Key: cpu\|overlap_summation |
| CPU\|Run | Time the VM is scheduled to run.<br>Key: cpu\|run_summation |
| CPU\|Entitlement Latest | Entitlement Latest.<br>Key: cpu\|entitlement_latest |
| CPU\|Total Capacity (MHz) | Total CPU capacity allocated to the virtual machine.<br>Key: cpu\|vm_capacity_provisioned |
| CPU\|Peak vCPU Ready | The highest CPU Ready among the virtual CPUs.<br>Key: cpu\|peak_vcpu_ready |
| CPU\|Peak vCPU Usage | The highest CPU Usage among the virtual CPU, compared with the static configured CPU frequency. A constantly high number indicates that one or more of the CPUs have high utilization.<br>Key: cpu\|peak_vcpu_usage |
| CPU\|20-second Peak CPU System (%) | The highest system CPU, measured as a peak of any 20-second average during the collection interval.<br>Key: cpu\|20-second peak cpu system |
| CPU\|20-second Peak vCPU Co-Stop (%) | The highest CPU Co-Stop among any of the vCPU, measured as a peak of any 20-second average during the collection interval.<br>Key: cpu\|20-second peak vcpu co-stop |
| CPU\|20-second Peak vCPU IO-Wait(%) | The highest CPU Other Wait among any of the vCPU, measured as a peak of any 20-second average during the collection interval. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: cpu\|20-second peak vcpu io-wait |
| CPU\|20-second Peak vCPU Overlap (ms) | The highest CPU Overlap among any of the vCPU, measured as a peak of any 20-second average during the collection interval.<br>Key: cpu\|20-second peak vcpu overlap |
| CPU\|20-second Peak vCPU Ready (%) | The highest CPU Ready among any of the vCPU, measured as a peak of any 20-second average during the collection interval.<br>Key: cpu\|20-second peak vcpu ready |
| CPU\|20-second Peak vCPU Swap Wait (%) | The highest CPU Swap Wait among any of the vCPU, measured as a peak of any 20-second average during the collection interval.<br>Key: cpu\|20-second peak vcpu swap wait |
| CPU \| vCPU Usage Disparity | The absolute gap between the highest vCPU Usage and the lowest vCPU Usage.<br>Key: cpu\|vcpu_usage_disparity |

**CPU Utilization for Resources Metrics for Virtual Machines**

CPU utilization for resources metrics provides information about resource CPU use.

| Metric Name | Description |
|---|---|
| rescpu\|CPU Active (%) (*interval*) | The average active time (actav) or peak active time (actpk) for the CPU during various intervals.<br>Key:<br>rescpu\|actav1_latest<br>rescpu\|actav5_latest<br>rescpu\|actav15_latest<br>rescpu\|actpk1_latest<br>rescpu\|actpk5_latest<br>rescpu\|actpk15_latest |
| rescpu\|CPU Running (%) (*interval*) | The average runtime (runav) or peak active time (runpk) for the CPU during various intervals.<br>Key:<br>rescpu\|runav1_latest<br>rescpu\|runav5_latest<br>rescpu\|runav15_latest<br>rescpu\|runpk1_latest |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | rescpu\|runpk5_latest<br><br>rescpu\|runpk15_latest |
| rescpu\|CPU Throttled (%) (*interval*) | Amount of CPU resources over the limit that were refused, average over various intervals.<br>Key:<br><br>rescpu\|maxLimited1_latest<br>rescpu\|maxLimited5_latest<br><br>rescpu\|maxLimited15_latest |
| rescpu\|Group CPU Sample Count | The sample CPU count.<br>Key: rescpu\|sampleCount_latest |
| rscpu\|Group CPU Sample Period (ms) | The sample period.<br>Key: rscpu\|samplePeriod_latest |

**Memory Metrics for Virtual Machines**

Memory metrics provide information about memory use and allocation.

| Metric Name | Description |
|---|---|
| Mem\|Host Active (KB) | Host active memory use in kilobytes.<br>Key: mem\|host_active |
| Mem\|Contention (KB) | Memory contention in kilobytes.<br>Key: mem\|host_contention |
| Mem\|Contention (%) | Percent memory contention.<br>Key: mem\|host_contentionPct |
| Mem\|Guest Configured Memory (KB) | Guest operating system configured memory in kilobytes.<br>Key: mem\|guest_provisioned |
| Mem\|Guest Active Memory (%) | Percent guest operating system active memory.<br>Key: mem\|guest_activePct |
| Mem\|Guest Non-Pageable Memory (KB) | Guest operating system non-pageable memory in kilobytes.<br>Key: mem\|guest_nonpageable_estimate |
| Mem\|Reservation Used | Memory Reservation Used.<br>Key: mem\|reservation_used |
| Mem\|Effective Limit | Memory Effective Limit.<br>Key: mem\|effective_limit |
| Mem\|Demand for aggregation | Host demand for aggregation.<br>Key: mem\|host_demand_for_aggregation |
| Mem\|Balloon (%) | Percentage of total memory that has been reclaimed via ballooning.<br>Key: mem\|balloonPct |
| Mem\|Guest Usage (KB) | This metric shows the amount of memory the VM uses. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: mem\|guest_usage |
| Mem\|Guest Demand (KB) | Guest operating system demand in kilobytes.<br>Key: mem\|guest_demand |
| Mem\|Guest Non-Pageable Memory (KB) | Guest operating system non-pageable memory in kilobytes.<br>Key: mem\|host_nonpageable_estimate |
| Mem\|Host Demand (KB) | Memory demand in kilobytes.<br>Key mem\|host_demand |
| Mem\|Host Workload | Host Workload (%).<br>Key: host_workload |
| Mem\|Zero (KB) | Amount of memory that is all 0.<br>Key: mem\|zero_average |
| Mem\|Swapped (KB) | This metric shows how much memory is being swapped. Meaning, the amount of unreserved memory in kilobytes.<br>Key: mem\|swapped_average |
| Mem\|Swap Target (KB) | Amount of memory that can be swapped in kilobytes.<br>Key: mem\|swaptarget_average |
| Mem\|Swap In (KB) | Swap-in memory in kilobytes.<br>Key: mem\|swapin_average |
| Mem\|Balloon Target (KB) | Amount of memory that can be used by the virtual machine memory control.<br>Key: mem\|vmmemctltarget_average |
| Mem\|Consumed (KB) | Amount of host memory consumed by the virtual machine for guest memory in kilobytes.<br>Key: mem\|consumed_average |
| Mem\|Overhead (KB) | Memory overhead in kilobytes.<br>Key: mem\|overhead_average |
| Mem\|Swap In Rate (KBps) | Rate at which memory is swapped from disk into active memory during the interval.<br>Key: mem\|swapinRate_average |
| Mem\|Active Write (KB) | Active writes in kilobytes.<br>Key: mem\|activewrite_average |
| Mem\|Compressed (KB) | Compressed memory in kilobytes.<br>Key: mem\|compressed_average |
| Mem\|Compression Rate (KBps) | Compression rate in kilobytes per second.<br>Key: mem\|commpressionRate_average |
| Mem\|Decompression Rate (KBps) | Decompression rate in kilobytes per second.<br>Key: mem\|decompressionRate_average |
| Mem\|Overhead Max (KB) | Maximum overhead in kilobytes.<br>Key: mem\|overheadMax_average |
| Mem\|Zip Saved (KB) | Zip-saved memory in kilobytes.<br>Key: mem\|zipSaved_latest |
| Mem\|Zipped (KB) | Zipped memory in kilobytes.<br>Key: mem\|zipped_latest |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Mem\|Entitlement | Amount of host physical memory the VM is entitled to, as determined by the ESX schedule.<br>Key: mem\|entitlement_average |
| Mem\|Capacity Contention | Capacity Contention.<br>Key: mem\|capacity.contention_average |
| Mem\|Swap In Rate from Host Cache | Rate at which memory is being swapped from host cache into active memory.<br>Key: mem\|llSwapInRate_average |
| Mem\|Swap Out Rate to Host Cache | Rate at which memory is being swapped to host cache from active memory.<br>Key: mem\|llSwapOutRate_average |
| Mem\|Swap Space Used in Host Cache | Space used for caching swapped pages in the host cache.<br>Key: mem\|llSwapUsed_average |
| Mem\|Overhead Touched | Actively touched overhead memory (KB) reserved for use as the virtualization overhead for the VM.<br>Key: mem\|overheadTouched_average |
| Memory\|VM Memory Demand (kb) | Key: mem\|vmMemoryDemand |
| Memory\|Consumed (%) | Key: mem\|consumedPct |
| Mem\|Utilization (KB) | Memory used by the virtual machine. Reflects the guest OS memory required for vSphere and certain VMTools versions or for virtual machine consumption.<br>Key: mem\|vmMemoryDemand |
| Mem\|Total Capacity (KB) | Memory resources allocated to powered on virtual machine.<br>Key: mem\|guest_provisioned |
| Mem\|20-second Peak Contention (%) | The highest Memory Contention, measured as peak of any 20-second average during the collection interval.<br><br>Key: guest\|20-second_peak_contention |
| Guest\|Peak Guest OS Page-out/rate within collection cycle | This metric shows the highest memory page-out rate reported by guest operating system, measured as peak of any 20-second average during the collection interval.<br>Key: guest\|20_sec_peak_page.outRate_latest |
| Guest\|Needed Memory | Amount of memory needed for the Guest OS to perform optimally. This memory is considered as a cache for the disk and is a little more than the actual used memory.<br><br>Key: guest\|mem.needed_latest |
| Guest\|Free Memory | Amount of memory that is not used but is readily available. If the cache is high, a low free memory does not mean that the Guest OS needs more memory.<br><br>Key: guest\|mem.free_latest |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Guest\|Physical Usable Memory | Amount of memory available to the Guest OS. Meaning, this amount is close to the amount of configured memory to the VM.<br><br>Key: guest\|mem.physUsable_latest |
| Guest\|20-second Peak Disk Queue Length | The highest Disk Queue Length, measured as peak of any 20-second average during the collection interval.<br><br>Key: guest\|20-second_peak_disk_queue_length |
| Guest\|20-second Peak Run Queue | The highest Run Queue, measured as peak of any 20-second average during the collection interval.<br><br>Key: guest\|20-second_peak_run_queue |
| Guest\|20-second Peak CPU Context Switch Rate | The highest CPU Context Switch Rate, measured as peak of any 20-second average during the collection interval.<br><br>Key: guest\|20-second_peak_cpu_context switch rate |

**Datastore Metrics for Virtual Machines**

Datastore metrics provide information about datastore use.

| Metric Name | Description |
|---|---|
| Datastore\|Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: datastore\|commandsAveraged_average |
| Datastore\|Outstanding IO requests | OIO for datastore.<br>Key: datastore\|demand_oio |
| Datastore\|Number of Outstanding IO Operations | Number of outstanding IO operations.<br>Key: datastore\|oio |
| Datastore\|Demand | Datastore demand.<br>Key: datastore\|demand |
| Datastore\|Total Latency (ms) | The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency.<br>Key: datastore\|totalLatency_average |
| Datastore\|Total Throughput (KBps) | Usage Average (KBps).<br>Key: datastore\|usage_average |
| Datastore\|Used Space (MB) | Used space in megabytes.<br>Key: datastore\|used |
| Datastore\|Not Shared (GB) | Space used by VMs that is not shared.<br>Key: datastore\|notshared |
| Datastore\|Read IOPS | Average number of read commands issued per second during the collection interval. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: datastore\|numberReadAveraged_average |
| Datastore\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: datastore\|numberWriteAveraged_average |
| Datastore\|Read Throughput (KBps) | This metric shows the amount of data that the VM reads to the datastore per second.<br>Key: datastore\|read_average |
| Datastore\|Read Latency (ms) | Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency.<br>Key: datastore\|totalReadLatency_average |
| Datastore\|Write Latency (ms) | Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency.<br>Key: datastore\|totalWriteLatency_average |
| Datastore\|Write Throughput (KBps) | This metric shows the amount of data that the VM writes to the datastore per second.<br>Key: datastore\|write_average |
| Datastore\|Highest Latency | Highest Latency.<br>Key: datastore\|maxTotalLatency_latest |
| Datastore\|Total Latency Max | Total Latency Max (ms).<br>Key: datastore\|totalLatency_max |

**Disk Metrics for Virtual Machines**

Disk metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| Disk Space\|vSAN Overhead (GB) | Displays the extra virtual machine disk space used by the vSAN system.<br>Disk Space\|vSAN Overhead (GB) |
| Disk Space\|VMUsedWithoutOverhead(GB) | Displays the virtual machine disk space without vSAN overhead.<br>Disk Space\|Usage Without Overhead (GB) |
| Disk Space<br>**NOTE**<br>The Disk Space metrics are displayed by the virtual machine object only when the vSAN adapter is configured with vCenter. | |
| Disk\|Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: disk\|numberReadAveraged_average |
| Disk\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: disk\|numberWriteAveraged_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Disk\|Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: disk\|commandsAveraged_average |
| Disk\|Total Throughput (KBps) | Use rate in kilobytes per second.<br>Key: disk\|usage_average |
| Disk\|I/O Usage Capacity | This metric is a function of storage\|usage_average and disk\|workload. Storage\|usage_average is an average over all storage devices. This means that disk\|usage_capacity is not specific to the selected VM or the host of the VM.<br>Key: disk\|usage_capacity |
| Disk\|Number of Outstanding IO Operations | Number of outstanding IO operations.<br>Key: disk\|diskoio |
| Disk\|Queued Operations | Queued operations.<br>Key: disk\|diskqueued |
| Disk\|Demand (%) | Percent demand.<br>Key: disk\|diskdemand |
| Disk \|Total Queued Outstanding Operations | Sum of Queued Operation and Outstanding Operations.<br>Key: disk \|sum_queued_oio |
| Disk\|Max Observed OIO | Max Observed IO for a disk.<br>Key: disk\|max_observed |
| Disk\|Read Throughput KBps) | Amount of data read in the performance interval.<br>Key: disk\|read_average |
| Disk\|Write Throughput (KBps) | Amount of data written to disk in the performance interval.<br>Key: disk\|write_average |
| Disk\|Bus Resets | The number of bus resets in the performance interval.<br>Key: disk\|busResets_summation |
| Disk\|Commands canceled | The number of disk commands canceled in the performance interval.<br>Key: disk\|commandsAborted_summation |
| Disk\|Highest Latency | Highest latency.<br>Key: disk\|maxTotalLatency_latest |
| Disk\|SCSI Reservation Conflicts | SCSI Reservation Conflicts.<br>Key: disk\|scsiReservationConflicts_summation |
| Disk\|Read Latency (ms) | The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency.<br>Key: disk\|totalReadLatency_average |
| Disk\|Write Latency (ms) | The average amount of time taken for a write from the perspective of a Guest OS. This is the sum of Kernel Write Latency and Physical Device Write Latency.<br>Key: disk\|totalWriteLatency_average |
| Disk\|Total Latency (ms) | The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Kernel Command Latency and Physical Device Command Latency.<br>Key: disk\|totalLatency_average |

**Virtual Disk Metrics for Virtual Machines**

Virtual disk metrics provide information about virtual disk use.

| Metric Name | Description |
|---|---|
| Virtual Disk:<scsi_controller>\|IOPS per GB | Displays the disk IO per second per Gigabyte of storage.<br>Key: virtualDisk:<scsi_controller>\|iopsPerGB |
| Virtual Disk\|Total Throughput | Amount of data read from/written to storage in a second. This is averaged over the reporting period.<br>Key: virtualDisk\|usage |
| VirtualDisk\|Total Latency | Total latency.<br>Key: virtualDisk\|totalLatency |
| VirtualDisk\|Total IOPS | Average number of commands per second.<br>Key: virtualDisk\|commandsAveraged_average |
| VirtualDisk\|Read Requests | Average number of read commands issued per second to the virtual disk during the collection interval.<br>Key: virtualDisk\|numberReadAveraged_average |
| VirtualDisk\|Write Requests | Average number of write commands issued per second to the virtual disk during the collection interval.<br>Key: virtualDisk\|numberWriteAveraged_average |
| VirtualDisk\|Read Throughput (KBps) | Rate of reading data from the virtual disk in kilobytes per second.<br>Key: virtualDisk\|read_average |
| VirtualDisk\|Read Latency (ms) | Average amount of time for a read operation from the virtual disk. Total latency = kernel latency + device latency.<br>Key: virtualDisk\|totalReadLatency_average |
| VirtualDisk\|Write Latency (ms) | Average amount of time for a write operation to the virtual disk. Total latency = kernel latency + device latency.<br>Key: virtualDisk\|totalWriteLatency_average |
| VirtualDisk\|Write Throughput (KBps) | Rate of writing data from the virtual disk in kilobytes per second.<br>Key: virtualDisk\|write_average |
| VirtualDisk\|Bus Resets | The number of bus resets in the performance interval.<br>Key: virtualDisk\|busResets_summation |
| VirtualDisk\|Commands Aborted | The number of disk commands canceled in the performance interval.<br>Key: virtualDisk\|commandsAborted_summation |
| VirtualDisk\|Read Load | Storage DRS virtual disk metric read load. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: virtualDisk\|readLoadMetric_latest |
| VirtualDisk\|Outstanding Read Requests | Average number of outstanding read requests to the virtual disk.<br>Key: virtualDisk\|readOIO_latest |
| VirtualDisk\|Write Load | Storage DRS virtual disk write load.<br>Key: virtualDisk\|writeLoadMetric_latest |
| VirtualDisk\|Outstanding Write Requests | Average number of outstanding write requests to the virtual disk.<br>Key: virtualDisk\|writeOIO_latest |
| VirtualDisk\|Number of Small Seeks | Small Seeks.<br>Key: virtualDisk\|smallSeeks_latest |
| VirtualDisk\|Number of Medium Seeks | Medium Seeks.<br>Key: virtualDisk\|mediumSeeks_latest |
| VirtualDisk\|Number of Large Seeks | Large Seeks.<br>Key: virtualDisk\|largeSeeks_latest |
| VirtualDisk\|Read Latency (microseconds) | Read Latency in microseconds.<br>Key: virtualDisk\|readLatencyUS_latest |
| VirtualDisk\|Write Latency (microseconds) | Write Latency in microseconds.<br>Key: virtualDisk\|writeLatencyUS_latest |
| VirtualDisk\|Average Read request size | Read IO size.<br>Key: virtualDisk\|readIOSize_latest |
| VirtualDisk\|Average Write request size | Write IO size.<br>Key: virtualDisk\|writeIOSize_latest |
| Virtual Disk\|Outstanding IO requests (OIOs) | Key: virtualDisk\|vDiskOIO |
| Virtual Disk\|Used Disk Space (GB) | Key: virtualDisk\|actualUsage |
| Virtual Disk\|Peak Virtual Disk IOPS | The highest disk IO per second among the virtual disks. A constantly high number indicates that one or more virtual disks are sustaining high IOPS.<br>Key: virtualDisk\|peak_vDisk_iops |
| Virtual Disk\|Peak Virtual Disk Read Latency | The highest read latency among the virtual disks. A high number indicates that one or more virtual disks are experiencing poor performance.<br>Key: virtualDisk\|peak_vDisk_readLatency |
| Virtual Disk\|Peak Virtual Disk Write Latency | The highest write latency among the virtual disks. A high number indicates that one or more virtual disks are experiencing poor performance.<br>Key: virtualDisk\|peak_vDisk_writeLatency |
| Virtual Disk\|20-secod Peak Latency (ms) | The highest latency among any of the virtual disk, measured as peak of any 20-second average during the collection interval.<br>Key: virtualDisk\|20-second_peak_latency |
| Virtual Disk\|Peak Virtual Disk throughput | The highest disk throughput among the virtual disks.<br>Key: virtualDisk\|peak_vDisk_throughpu |

## Guest File System Metrics for Virtual Machines

Guest file system metrics provide information about guest file system capacity and free space.

The data for these metrics is only displayed when VMware Tools has been installed on the virtual machines. If VMware Tools is not installed, features dependent on these metrics, including capacity planning for virtual machine guest storage, will not be available.

| Metric Name | Description |
|---|---|
| Guest file system\|Guest File System Capacity (MB) | Total capacity on guest file system in megabytes.<br>Key: guestfilesystem\|capacity |
| Guest file system\|Guest File System Free (MB) | Total free space on guest file system in megabytes.<br>Key: guestfilesystem\|freespace |
| Guest file system\|Guest File System Usage (%) | Percent guest file system.<br>Key: guestfilesystem\|percentage |
| Guest file system\|Guest File System Usage | Total usage of guest file system.<br>From VMware Cloud Foundation Operations 6.7 and onwards, this metric is measured in GBs.<br><br>Key: guestfilesystem\|usage |
| Guest file system\|Total Guest File System Capacity (GB) | This metric displays the amount of disk space allocated for the VM.<br>Correlate other metrics with this metric to indicate if changes occur in the disk space allocation for the VM.<br><br>Key: guestfilesystem\|capacity_total |
| Guest file system\|Total Guest File System Usage (%) | This metric displays the amount of display space being used out of the total allocated disk space.<br>Use his metric to track if the overall usage is stable, or if it reaches the limits. Do not include VMs with a disk space usage of >95% since this might impact your system.<br><br>Key: guestfilesystem\|percentage_total |
| Guest file system\|Total Guest File System Usage | Total usage of guest file system.<br>Key: guestfilesystem\|usage_total |
| Guest file system\|Utilization (GB) | Storage space used by the Guest OS file systems. The disk space is available only if VM tools are installed and running. If the VM tools are not installed, the disk space capacity is not applicable.<br>Key: guestfilesystem\|usage_total |
| Guest file system\|Total Capacity (GB) | Storage space used by the Guest OS file systems. The disk space is available only if VM tools are installed and running. If the VM tools are not installed, the disk space capacity is not applicable.<br>Key: guestfilesystem\|capacity_total |

**Network Metrics for Virtual Machines**

Network metrics provide information about network performance.

| Metric Name | Description |
|---|---|
| Network\|Peak Network Packet per second within collection cycle | This metric shows the highest VM packets per second rate, measured as peak of any 20-second average during the collection interval.<br>Key: net\|20_sec_peak_packetsPerSec |
| Net\|Total Throughput (KBps) | The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.<br>Key: net\|usage_average |
| Net\|Data Transmit Rate (KBps) | This metric shows the rate of data being sent by the VM per second.<br>Key: net\|transmitted_average |
| Net\|Data Receive Rate (KBps) | This metric shows the rate of data received by the VM per second.<br>Key: net\|received_average |
| Net\|Packets per second | Number of packets transmitted and received per second.<br>Key: net\|PacketsPerSec |
| Net\|Packets Received | Number of packets received in the performance interval.<br>Key: net\|packetsRx_summation |
| Net\|Packets Transmitted | Number of packets transmitted in the performance interval.<br>Key: net\|packetsTx_summation |
| Net\|Transmitted Packets Dropped | This metric shows the number of transmitted packets dropped in the collection interval<br>Key: net\|dropppedTx_summation |
| Net\|Packets Dropped (%) | Percentage of packets dropped.<br>Key: net\|droppedPct |
| Net\|Packets Dropped | Number of packets dropped in the performance interval.<br>Key: net\|dropped |
| Net\|Broadcast Packets Transmitted | Number of broadcast packets transmitted during the sampling interval.<br>Key: net\|broadcastTx_summation |
| Net\|Broadcast Packets Received | Number of broadcast packets received during the sampling interval.<br>Key: net\|broadcastRx_summation |
| Net\|Multicast Packets Received | Number of multicast packets received.<br>Key: net\|multicastRx_summation |
| Net\|Multicast Packets Transmitted | Number of multicast packets transmitted.<br>Key: net\|multicastTx_summation |
| Net\|VM to Host Data Transmit Rate | Average amount of data transmitted per second between VM and host.<br>Key: net\|host_transmitted_average |
| Net\|VM to Host Data Receive Rate | Average amount of data received per second between VM and host.<br>Key: net\|host_received_average |
| Net\|VM to Host Usage Rate | The sum of the data transmitted and received for all the NIC instances between VM and host.<br>Key: net\|host_usage_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Net\|20-second Peak Usage Rate (KBps) | The higest Usage Rate, measured as peak of any 20 second average during the collection interval.<br>Key: net\|20-second_peak_usage_rate |

**System Metrics for Virtual Machines**

System metrics for virtual machines provide general information about the virtual machine, such as its build number and running state.

| Metric Name | Description |
|---|---|
| Sys\|Powered ON | Powered on virtual machines. 1 if powered on, 0 if powered off, -1 if unknown<br>Key: sys\|poweredOn |
| Sys\|OS Uptime | Total time elapsed, in seconds, since last operating system start.<br>Key: sys\|osUptime_latest |

**Power Metrics for Virtual Machines**

Power metrics provide information about power use.

| Metric Name | Description |
|---|---|
| Power\|Total Energy Consumed in the collection period (Wh) | Displays the total electricity consumed based on the time interval selected. The default collection cycle is set to 5 mins. You can continue using the default setting or edit it for each adapter instance. For example, if the time interval is set to its default value, the value represents the energy consumed per 5 mins. |
| Power\|Total Power Consumed By VM (Wh) | Displays the total power consumed by a virtual machine in an hour. The data collected is over a period of an hour and published along with the other metrics in VMware Aria OperationsVMware Cloud Foundation Operations. In case of a connectivity or availability issue in VMware Aria OperationsVMware Cloud Foundation Operations or vCenter adapter instance, this hourly metric might not be published and the missed value during this period does not get recalculated. Once the connection is re-established, the next data points get published. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | **NOTE**<br>This metric is deactivated by default. You can activate it from the Policies page. For more information, see Metrics and Properties Details in the *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide*. This metric is deactivated by default. You can activate it from the Policies page. For more information, see Metrics and Properties Details. |
| Power\|Power (Watt) | Average power use in watts. |
| Power\|Current Power Consumption Rate (Watt) | The power consumption rate per second, averaged over the reporting period.<br>Key: power\|power_average |
| Power\|(DEP) Energy (Joule) | Total energy consumed in joules.<br>Key: power\|energy_summation |

## Disk Space Metrics for Virtual Machines

Disk space metrics provide information about disk space use.

| Metric Name | Description |
|---|---|
| Diskspace\|Provisioned Space (GB) | Provisioned space in gigabytes.<br>Key: diskspace\|provisioned |
| Diskspace\|Provisioned Space for VM | Provisioned space for VM.<br>Key: diskspace\|provisionedSpace |
| Diskspace\|Snapshot Space (GB) | Space used by snapshots.<br>Key: diskspace\|snapshot |
| Diskspace\|Virtual machine used (GB) | Space used by virtual machine files in gigabytes.<br>Key: diskspace\|perDsUsed |
| Diskspace\|Active not shared | Unshared disk space used by VMs excluding snapshot.<br>Key: diskspace\|activeNotShared |

## Storage Metrics for Virtual Machines

Storage metrics provide information about storage use.

| Metric Name | Description |
|---|---|
| Storage\|Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: storage\|commandsAveraged_average |
| Storage\|Contention (%) | Percent contention.<br>Key: storage\|contention |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Storage\|Read Throughput (KBps) | Read throughput rate in kilobytes per second.<br>Key: storage\|read_average |
| Storage\|Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: storage\|numberReadAveraged_average |
| Storage\|Total Latency (ms) | Total latency in milliseconds.<br>Key: storage\|totalLatency_average |
| Storage\|Total Usage (KBps) | Total throughput rate in kilobytes per second.<br>Key: storage\|usage_average |
| Storage\|Write Throughput (KBps) | Write throughput rate in kilobytes per second.<br>Key: storage\|write_average |
| Storage\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: storage\|numberWriteAveraged_average |

**Summary Metrics for Virtual Machines**

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| Summary\|Availability % | Displays the uptime of Guest OS, expressed as a percentage of the collection period.<br>Key: summary\|availability_kpi |
| Summary\|Running | Number of running virtual machines.<br>Key: summary\|running |
| Summary\|Desktop Status | Horizon view desktop status.<br>Key: summary\|desktop_status |
| Summary\|Configuration\|Type | Indicates the type of virtual machine object based on which you can identify the type of virtual machine. The valid values for the virtual machine object property are:<br>• default - represents a regular virtual machine<br>• template - represents a powered off virtual machine template.<br>• srm_placeholder - represents a powered on Site Recovery Manager virtual machine.<br>• ft_primary - represents the primary Fault Tolerance virtual machine.<br>• ft_secondary - represents the secondary Fault Tolerance virtual machine.<br>Key: summary\|config\|type |
| Summary\|Guest Operating System\|Guest OS Full Name | Displays the guest operating system name.<br>Key: summary\|guest os full name |
| Summary\|Oversized\|Potential Memory | Displays the oversized potential memory.<br>Key: summary\|oversized\|potentialMemConsumed |
| Summary\|Undersized\|Potential CPU Usage | Displays the undersized potential CPU used. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: summary\|undersized\|potentialCpuUsage |
| Summary\|Undersized\|Potential Memory | Displays the undersized potential memory used.<br>Key: summary\|undersized\|potentialMemUsage |
| Reclaimable Idle | Boolean flag indicating whether VM is considered as reclaimable because it is in Idle state.<br>Key: summary\|idle |
| Reclaimable Powered Off | Boolean flag indicating whether VM is considered as reclaimable because it is in powered off state.<br>Key: summary\| poweredOff |
| Reclaimable Snapshot Space (GB) | Reclaimable snapshot space.<br>Key: summary\| snapshotSpace |

**Cost Metrics for Virtual Machines**

Cost metrics provide information about the cost.

| Metric Name | Description |
|---|---|
| Monthly OS Labor Cost | Monthly operating system labor cost of the virtual machine.<br>Key: cost\|osLaborTotalCost |
| Monthly Projected Total Cost | Virtual machine cost projected for full month.<br>Key: Cost\|monthlyProjectedCost |
| Monthly VI Labor Cost | Monthly virtual infrastructure labor cost of the virtual machine.<br>Key: cost\|viLaborTotalCost |
| MTD Compute Total Cost | Total compute cost (including CPU and memory) of the virtual machine.<br>Key: cost\|compTotalCost |
| MTD CPU Cost | Month to Date Virtual Machine CPU Cost. It is based on utilization. The more the VM uses, the higher its cost.<br>Key: cost\|cpuCost |
| MTD Monthly Cost | Month to date direct cost (comprising of OS labor, VI labor and any windows desktop instance license) of the virtual machine. It also comprises of the additional and application cost of the virtual machine.<br>Key: cost\|vmDirectCost |
| MTD Memory Cost | Month to Date Memory Cost of Virtual Machine. It is based on utilization. The more the VM uses, the higher its cost.<br>Key: cost\|memoryCost |
| MTD Storage Cost | Month to date storage cost of the virtual machine.<br>Key: cost\|storageCost |
| MTD Total Cost | Month to date total compute cost (including CPU and memory) of the virtual machine.<br>Key: cost\|monthlyTotalCost |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Potential Savings | Reclaimable cost of VM for being either idle, powered-off, or having snapshots.<br>Key: cost\|reclaimableCost |
| Cost\|Allocation\|MTD VM CPU Cost (Currency) | Month to Date Virtual Machine CPU Cost computed based on resource overcommit ratio set for its parent cluster in policy.<br>cost\|allocation\|allocationBasedCpuMTDCost |
| Cost\|Allocation\|MTD VM Memory Cost (Currency) | Month to Date Virtual Machine CPU Memory cost computed based on resource overcommit ratio set for its parent cluster in policy.<br>cost\|allocation\|allocationBasedMemoryMTDCost |
| Cost\|Allocation\|MTD VM Storage Cost (Currency) | Month to Date Virtual Machine CPU Storage cost computed based on resource overcommit ratio set for its parent cluster (or datastore cluster) in policy.<br>cost\|allocation\|allocationBasedStorageMTDCost |
| Cost\|Allocation\|MTD VM Total Cost (Currency) | Month to Date Virtual Machine Total Cost is the summation of the CPU Cost, Memory Cost, Storage Cost and Direct Cost, based on overcommit ratios set in policy for the parent cluster or datastore cluster.<br>cost\|allocation\|allocationBasedTotalCost |
| Cost\|Effective Daily Cpu Cost (Currency) | Daily CPU cost of the selected virtual machine. |
| Cost\|Effective Daily Memory Cost (Currency) | Daily Memory cost of the selected virtual machine. |
| Cost\|Effective Daily Storage Cost (Currency) | Daily Storage cost of the selected virtual machine. |
| Cost\|Daily Additional Cost | Daily Additional cost of the selected virtual machine. |
| Cost\|Effective Daily Cost (Currency) | Effective Daily cost is the sum of effective daily CPU cost + effective daily memory cost + effective daily storage cost + daily additional cost. |
| Cost\|Effective MTD Cost (Currency) | Effective MTD cost is the sum of effective daily CPU cost from beginning of month until now + effective daily memory cost from beginning of month until now + effective daily storage cost from beginning of month until now + daily additional cost from beginning of month until now. |

**Virtual Hardware Metrics for Virtual Machines**

| Metric Name | Description |
|---|---|
| Configuration\|Hardware\|Number of CPU cores per socket | This metric displays the number of CPU cores per socket. |
| Configuration\|Hardware\|Number of virtual CPUs | This metric displays the number of CPUs in the virtual machine. |
| Configuration\|Hardware\|Number of virtual sockets: | This metric displays the number of virtual sockets in the virtual machine. |
| Configuration\|Hardware\|Memory: | This metric displays the memory used in the virtual machine. |
| Configuration\|CPU Resource Allocation\|Limit | This metric displays the resource allocation limit of the virtual machine. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Configuration\|CPU Resource Allocation\|Reservation | This metric displays the reserved resources for the virtual machine. |
| Configuration\|CPU Resource Allocation\|Shares\| | This metric displays the shared resources for the virtual machine. |
| Summary\|Guest Operating System\|Tools Version | This metric displays the tools version of the guest operating system. |
| Summary\|Guest Operating System\|Tools Version Status | This metric displays the status of the tools in the guest operating system. |
| Summary\|Guest Operating System\|Tools Running Status | This metric displays whether the tools are functional in the guest operating system. |
| Guest File System:/boot\|Partition Capacity (GB) | This metric displays the boot partition capacity in the guest file system. |
| Guest File System:/boot\|Partition Utilization (%) | This metric displays the boot partition usage percentage in the guest file system. |
| Guest File System:/boot\|Partition Utilization (GB) | This metric displays the boot partition used in the guest file system. |
| Virtual Disk\|Configured | This metric displays the disk space of the configured virtual disk. |
| Virtual Disk\|Label | This metric displays the disk label of the configured virtual disk. |
| Disk Space\|Snapshot Space | This metric displays the snap shot details of the virtual machine. |
| Network\|IP Address | This metric displays the IP address of the virtual machine. |
| Network\|MAC Address | This metric displays the MAC address of the virtual machine. |

**Disabled Instanced Metrics**

The instance metrics created for the following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that these metrics collect data by default but all the instanced metrics created for these metrics, do not collect data by default.

| Metric Name |
|---|
| Configuration\|Hardware\|Number of virtual CPUs |
| CPU\|Ready (%) |
| CPU\|Usage (MHz) |
| Net\|Broadcast Packets Transmitted |
| Net\|Data Transmit Rate (KBps) |
| Net\|Data Receive Rate (KBps) |
| Net\|Multicast Packets Transmitted |
| Net\|Packets Dropped |
| Net\|Packets Dropped (%) |
| Net\|pnicByteRx_average |
| Net\|pnicByteTx_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name |
|---|
| Net\|Transmitted Packets Dropped |
| Net\|Usage Rate (KBps) |
| VirtualDisk\|Read IOPS |
| VirtualDisk\|Read Latency (ms) |
| VirtualDisk\|Read Throughput (KBps) |
| VirtualDisk\|Total IOPS |
| VirtualDisk\|Total Latency |
| VirtualDisk\|Total Throughput (KBps) |
| Virtual Disk\|Used Disk Space (GB) |
| VirtualDisk\|Write IOPS |
| VirtualDisk\|Write Latency (ms) |
| VirtualDisk\|Write Throughput (KBps) |
| Datastore\|Outstanding IO requests |
| Datastore\|Read IOPS |
| Datastore\|Read Latency (ms) |
| Datastore\|Read Throughput (KBps) |
| Datastore\|Total IOPS |
| Datastore\|Total Latency (ms) |
| Datastore\|Total Throughput (KBps) |
| Datastore\|Write IOPS |
| Datastore\|Write Latency (ms) |
| Datastore\|Write Throughput (KBps) |
| Disk\|Total IOPS |
| Disk\|Total Throughput (KBps) |
| Disk\|Read Throughput KBps) |
| Disk\|Write Throughput (KBps) |
| Diskspace\|Access Time (ms) |
| Diskspace\|Virtual machine used (GB) |

**Disabled Metrics**

The following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

You can enable these metrics in the Policy workspace.

| Metric Name | Description |
|---|---|
| CPU\|50% of Recommended number of vCPUs to Remove | This metric is superseded by the capacity engine. cpu\|numberToRemove50Pct |
| CPU\|Capacity entitlement (mhz) | cpu\|capacity_entitlement |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| CPU\|Co-stop (msec) | Use the Co-Stop (%) metric instead of this metric.<br>cpu\|costop_summation |
| CPU\|Demand Over Capacity (mhz) | cpu\|demandOverCapacity |
| CPU\|Demand Over Limit (mhz) | Use Contention (%) metric instead of this metric.<br>cpu\|demandOverLimit |
| CPU\|Dynamic entitlement | cpu\|dynamic_entitlement |
| CPU\|Estimated entitlement | cpu\|estimated_entitlement |
| CPU\|Idle (%) | cpu\|idlePct |
| CPU\|Idle (msec) | cpu\|idle_summation |
| CPU\|Other Wait (msec) | cpu\|otherwait |
| CPU\|Normalized Co-stop (%) | Use the Co-Stop (%) metric instead of this metric.<br>cpu\|perCpuCoStopPct |
| CPU\|Provisioned vCPU(s) (Cores) | cpu\|corecount_provisioned |
| CPU\|Ready (msec) | Choose the Use Ready (%) metric instead of this metric.<br>cpu\|ready_summation |
| CPU\|Recommended Size Reduction (%) | cpu\|sizePctReduction |
| CPU\|Swap Wait (msec) | cpu\|swapwait_summation |
| CPU\|Total Wait (msec) | cpu\|wait |
| CPU\|Used (msec) | cpu\|used_summation |
| CPU\|Wait (msec) | cpu\|wait_summation |
| Datastore I/O\|Max Observed Number of Outstanding IO Operations | datastore\|maxObserved_OIO |
| Datastore I/O\|Max Observed Read Rate (kbps) | datastore\|maxObserved_Read |
| Datastore I/O\|Max Observed Reads per second | datastore\|maxObserved_NumberRead |
| Datastore I/O\|Max Observed Write Rate (kbps) | datastore\|maxObserved_Write |
| Datastore I/O\|Max Observed Writes per second | datastore\|maxObserved_NumberWrite |
| Disk Space\|Not Shared (gb) | diskspace\|notshared |
| Disk Space\|Number of Virtual Disks | diskspace\|numvmdisk |
| Disk Space\|Shared Used (gb) | diskspace\|shared |
| Disk Space\|Total disk space used (gb) | diskspace\|total_usage |
| Disk Space\|Total disk space (gb) | diskspace\|total_capacity |
| Disk Space\|Virtual Disk Used (gb) | diskspace\|diskused |
| Guest File System stats\|Total Guest File System Free (gab) | guestfilesystem\|freespace_total |
| Guest\|Active File Cache Memory (kb) | guest\|mem.activeFileCache_latest |
| Guest\|Context Swap Rate per second | guest\|contextSwapRate_latest |
| Guest\|Huge Page Size (kb) | guest\|hugePage.size_latest |
| Guest\|Page Out Rate per second | guest\|page.outRate_latest |
| Guest\|Total Huge Pages | guest\|hugePage.total_latest |
| Memory\|50% of Reclaimable Memory Capacity (gb) | This metric is superseded by the capacity engine.<br>mem\|wasteValue50PctInGB |
| Memory\|Balloon (kb) | mem\|vmmemctl_average |
| Memory\|Demand Over Capacity | mem\|demandOverCapacity |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Memory\|Demand Over Limit | mem\|demandOverLimit |
| Memory\|Granted (kb) | mem\|granted_average |
| Memory\|Guest Active (kb) | mem\|active_average |
| Memory\|Guest Dynamic Entitlement (kb) | mem\|guest_dynamic_entitlement |
| Memory\|Guest Workload (%) | mem\|guest_workload |
| Memory\|Host Demand with Reservation (kb) | mem\|host_demand_reservation |
| Memory\|Host Dynamic Entitlement (kb) | mem\|host_dynamic_entitlement |
| Memory\|Host Usage (kb) | mem\|host_usage |
| Memory\|Host Workload (%) | mem\|host_workload |
| Memory\|Latency (%) | Use the Memory Contention (%) metric instead of this metric.<br>mem\|latency_average |
| Memory\|Recommended Size Reduction (%) | mem\|sizePctReduction |
| Memory\|Shared (kb) | mem\|shared_average |
| Memory\|Swap Out Rate (kbps) | mem\|swapoutRate_average |
| Memory\|Usage (%) | mem\|usage_average |
| Memory\|Estimated entitlement | mem\|estimated_entitlement |
| Network I/O\|Data Receive Demand Rate (kbps) | net\|receive_demand_average |
| Network I/O\|Data Transmit Demand Rate (kbps) | net\|transmit_demand_average |
| Network I/O\|VM to Host Data Receive Rate (kbps) | net\|host_received_average |
| Network I/O\|VM to Host Data Transmit Rate (kbps) | net\|host_transmitted_average |
| Network I/O\|VM to Host Max Observed Received Throughput (kbps) | net\|host_maxObserved_Rx_KBps |
| Network I/O\|VM to Host Max Observed Throughput (kbps) | net\|host_maxObserved_KBps |
| Network I/O\|VM to Host Max Observed Transmitted Throughput (kbps) | net\|host_maxObserved_Tx_KBps |
| Network I/O\|VM to Host Usage Rate (kbps) | net\|host_usage_average |
| Network\|bytesRx (kbps) | net\|bytesRx_average |
| Network\|bytesTx (kbps) | net\|bytesTx_average |
| Network\|Demand (%) | Use absolute numbers instead of this metric.<br>net\|demand |
| Network\|I/O Usage Capacity | net\|usage_capacity |
| Network\|Max Observed Received Throughput (kbps) | net\|maxObserved_Rx_KBps |
| Network\|Max Observed Throughput (kbps) | net\|maxObserved_KBps |
| Network\|Max Observed Transmitted Throughput (kbps) | net\|maxObserved_Tx_KBps |
| Network\|Packets Received per second | net\|packetsRxPerSec |
| Network\|Packets Transmitted per second | net\|packetsTxPerSec |
| Network\|Received Packets Dropped | net\|droppedRx_summation |
| Storage\|Demand (kbps) | storage\|demandKBps |
| Storage\|Read Latency (msec) | storage\|totalReadLatency_average |
| Storage\|Write Latency (msec) | storage\|totalWriteLatency_average |
| Summary\|CPU Shares | summary\|cpu_shares |
| Summary\|Memory Shares | summary\|mem_shares |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Summary\|Number of Datastores | summary\|number_datastore |
| Summary\|Number of Networks | summary\|number_network |
| Summary\|Workload Indicator | summary\|workload_indicator |
| System\|Build Number | sys\|build |
| System\|Heartbeat | sys\|heartbeat_summation |
| System\|Product String | sys\|productString |
| System\|Uptime (sec) | sys\|uptime_latest |
| System\|vMotion Enabled | vMotion should be enabled for all. It is not necessary to track all VMs every five minutes. sys\|vmotionEnabled |

### Host System Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects many metrics for host systems, including CPU use, datastore, disk, memory, network, storage, and summary metrics for host system objects.

Capacity metrics can be calculated for host system objects. See Capacity Analytics Generated Metrics.

### GPU Metrics

GPU metrics provide information about the GPU usage and performance.

| Metric Name | Description |
|---|---|
| **Metrics Aggregated at Host Level** | |
| GPU\|Compute Utilization (%) | The compute utilization percentage of a GPU. |
| GPU\|Memory Usage (%) | Memory currently in use as a percentage of total available memory. |
| GPU\|Memory Used (KB) | The amount of GPU memory used in kilobytes. |
| GPU\|Number of GPUs | Number of GPUs. |
| GPU\|Total Memory (KB) | Total memory in kilobytes. |
| **GPU Level Metrics** | |
| GPU\|<GPU-ID>\|Compute Utilization (%) | The compute utilization percentage of a GPU. |
| GPU\|<GPU-ID>\|Memory Usage (%) | Memory currently in use as a percentage of total available memory. |
| GPU\|<GPU-ID>\|Memory Used (KB) | The amount of GPU memory used in kilobytes. |
| GPU\|<GPU-ID>\|Memory Reserved (KB) | The amount of GPU memory reserved in kilobytes. |
| GPU\|<GPU-ID>\|Total Memory (KB) | Total memory in kilobytes. |
| GPU\|<GPU-ID>\|Temperature (Celsius | The temperature of a GPU in degrees celsius. |
| GPU\|<GPU-ID>\|Power Used (Watt) | The power used by a GPU in watts. |

### Host System Metrics for ROI Dashboard

Host system metrics provide information about cost saving across vCenters

| Metric Name | Description |
|---|---|
| Cost\|Monthly Additional Total Cost | This metric shows the total sum of additional cost across all the vCenters for an entire month.<br>Key: cost\|additionalTotalCost |

### Configuration Metrics for Host Systems

Configuration metrics provide information about host system configuration.

| Metric Name | Description |
|---|---|
| Configuration\|Hyperthreading\|Active | Displays the hyperthreading status of the host.<br>Key: configuration\|hypwerthreading\|active |
| Configuration\|Hyperthreading\|Available | Displays whether the hyperthreading option is available for this host.<br>Key: configuration\|hypwerthreading\|available |
| Configuration\|Storage Device\|Multipath Info\|Total number of Active Path | Displays the amount of active path information for the storage device<br>Key: configuration\|storagedevice\|multipathinfo\|total numberofActive path |
| Configuration\|Storage Device\|Total number of path | Displays the total number of path for the storage device.<br>Key: configuration\|storagedevice\|total number of path |
| Configuration\|Failover Hosts | Failover Hosts.<br>Key: configuration\|dasConfig\|admissionControlPolicy\|failoverHost |

### Hardware Metrics for Host Systems

Hardware metrics provide information about host system hardware.

| Metric Name | Description |
|---|---|
| Hardware\|Number of CPUs | Number of CPUs for a host.<br>Key: hardware\|cpuinfo\|num_CpuCores |
| Hardware\|ServiceTag | Displays the service tag of the host system.<br>Key: hardware\|servicetag |

### CPU Usage Metrics for Host Systems

CPU usage metrics provide information about CPU use.

| Metric Name | Description |
|---|---|
| CPU\|Capacity Usage (%) | Percent CPU capacity used.<br>Key: cpu\|capacity_usagepct_average |
| CPU\|Usage (%) | Average CPU usage as a percentage.<br>Key: cpu\|usage_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| CPU\|CPU Contention (%) | This metric indicates the percentage of time the virtual machines in the ESXi hosts are unable to run because they are contending for access to the physical CPU(s). This is the average number of all VMs. Naturally, the number will be lower than the highest number experienced by the worst hit VM (a VM that suffers the highest CPU contention).<br>Use this metric to verify if the host is able to serve all of its VMs well.<br><br>When using this metric, ensure the number is within your expectation. The metric is affected by several factors so you need to watch both relative numbers and absolute numbers. Relative means a drastic change in value. This indicates that the ESXi is unable to service its VMs.<br><br>Absolute means that the real value is high and should be checked. One factor that impacts the CPU contention metric is CPU Power Management. If CPU Power Management clocks down the CPU speed from 3 GHz to 2 GHz that reduction in speed is taken into consideration. This is because the VM is not running at full speed.<br><br>Key: cpu\|capacity_contentionPct |
| CPU\|Demand (%) | This metric shows the percentage of CPU resources all the VMs would use if there was no CPU contention or any CPU limits set.<br>It represents the average active CPU load for the past five minutes.<br><br>Keep the number of this metric below 100% if you set Power Management to Maximum.<br><br>Key: cpu\|demandPct |
| CPU\|Demand (MHz) | CPU demand in megahertz. CPU utilization level based on descendant Virtual Machines utilization. Includes limits and overhead to run Virtual Machines, but not reservations.<br>Key: cpu\|demandmhz |
| CPU\|IO Wait (ms) | IO wait time in milliseconds.<br>Key: cpu\|iowait |
| CPU\|Number of CPU Sockets | Number of CPU sockets.<br>Key: cpu\|numpackages |
| CPU\|Overall CPU Contention (ms) | Overall CPU contention in milliseconds.<br>Key: cpu\|capacity_contention |
| CPU\|Provisioned Capacity (MHz) | Capacity in MHz of the physical CPU cores.<br>Key: cpu\|capacity_provisioned |
| CPU\|Provisioned virtual CPUs | Provisioned virtual CPUs.<br>Key: cpu\|corecount_provisioned |
| CPU\|Total Wait | CPU time spent in idle state. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: cpu\|wait |
| CPU\|Demand | CPU demand.<br>Key: cpu\|demand_average |
| CPU\|Usage (MHz) | CPU use in megahertz.<br>Key: cpu\|usagemhz_average |
| CPU\|Reserved Capacity (MHz) | The sum of the reservation properties of the (immediate) children of the host's root resource pool.<br>Key: cpu\|reservedCapacity_average |
| CPU\|Total Capacity (MHz) | Total CPU capacity in megahertz. Amount of CPU resources configured on the ESXi hosts.<br>Key: cpu\|capacity_provisioned |
| CPU\|Overhead (KB) | Amount of CPU overhead.<br>Key: cpu\|overhead_average |
| CPU\|Demand without overhead | Value of demand excluding any overhead.<br>Key: cpu\|demand_without_overhead |
| CPU\|Core Utilization (%) | Percent core utilization.<br>Key: cpu\|coreUtilization_average |
| CPU\|Utilization(%) | Percent CPU utilization.<br>Key: cpu\|utilization_average |
| CPU\|Core Utilization (%) | Core Utilization.<br>Key: cpu\|coreUtilization_average |
| CPU\|Utilization (%) | Utilization.<br>Key: cpu\|utilization_average |
| CPU\|Co-stop (ms) | Time the VM is ready to run, but is unable to due to co-scheduling constraints.<br>Key: cpu\|costop_summation |
| CPU\|Latency (%) | Percentage of time the VM is unable to run because it is contending for access to the physical CPUs.<br>Key: cpu\|latency_average |
| CPU\|Ready (ms) | Time spent in ready state.<br>Key: cpu\|ready_summation |
| CPU\|Run (ms) | Time the virtual machine is scheduled to run.<br>Key: cpu\|run_summation |
| CPU\|Swap wait (ms) | Amount of time waiting for swap space.<br>Key: cpu\|swapwait_summation |
| CPU\|Wait (ms) | Total CPU time spent in wait state.<br>Key: cpu\|wait_summation |
| CPU\|Provisioned Capacity | Provisioned capacity (MHz).<br>Key: cpu\|vm_capacity_provisioned |
| CPU\|Active Host Load For Balance (Long Term) | Active Host Load For Balance (Long Term).<br>Key: cpu\|acvmWorkloadDisparityPcttive_longterm_load |
| CPU\|Active Host Load For Balance (Short Term) | Active Host Load For Balance (Short Term).<br>Key: cpu\|active_shortterm_load |
| CPU\| CPU Model | Displays the host's CPU model.<br>Key: cpu\|cpu model |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| CPU\|Peak CPU Core Usage | The highest CPU Usage among the CPU cores. A constantly high number indicates that one or more physical cores have high utilization.<br>Key: cpu\|peak_cpu_core_usage |

**CPU Utilization for Resources Metrics for Host Systems**

CPU utilization for resources metrics provide information about CPU activity.

| Metric Name | Description |
|---|---|
| Rescpu\|CPU Active (%) (*interval*) | Average active time for the CPU over the past minute, past five minutes, and at one-minute, five-minute, and 15-minute peak active times.<br>Key:<br>rescpu\|actav1_latest<br>rescpu\|actav5_latest<br>rescpu\|actav15_latest<br>rescpu\|actpk1_latest<br>rescpu\|actpk5_latest<br>rescpu\|actpk15_latest |
| Rescpu\|CPU Running (%) (*interval*) | Average run time for the CPU over the past minute, past five minutes, past 15 minutes, and at one-minute, five-minute, and 15-minute peak times.<br>Key:<br>rescpu\|runav1_latest<br>rescpu\|runav5_latest<br>rescpu\|runav15_latest<br>rescpu\|runpk1_latest<br>rescpu\|runpk5_latest<br>rescpu\|runpk15_latest |
| Rescpu\|CPU Throttled (%) (*interval*) | Scheduling limit over the past minute, past five minutes, and past 15 minutes.<br>Key:<br>rescpu\|maxLimited1_latest<br>rescpu\|maxLimited5_latest<br>rescpu\|maxLimited15_latest |
| Rescpu\|Group CPU Sample Count | Group CPU sample count. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: rescpu\|sampleCount_latest |
| Rescpu\|Group CPU Sample Period (ms) | Group CPU sample period in milliseconds.<br>Key: rescpu\|samplePeriod_latest |

**Datastore Metrics for Host Systems**

Datastore metrics provide information about datastore use.

| Metric Name | Description |
|---|---|
| Datastore\|Outstanding IO requests | OIO for datastore.<br>Key: datastore\|demand_oio |
| Datastore\|Commands Averaged | Average number of commands issued per second during the collection interval.<br>Key: datastore\|commandsAveraged_average |
| Datastore\|Number of Outstanding IO Operations | Number of outstanding IO operations.<br>Key: datastore\|oio |
| Datastore\|Total Latency (ms) | The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency.<br>Key: datastore\|totalLatency_average |
| Datastore\|Total Throughput (KBps) | Usage Average (KBps).<br>Key: datastore\|usage_average |
| Datastore\|Demand | Demand.<br>Key: datastore\|demand |
| Datastore\|Storage I/O Control aggregated IOPS | Aggregate number of IO operations on the datastore.<br>Key: datastore\|datastoreIops_average |
| Datastore\|Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: datastore\|numberReadAveraged_average |
| Datastore\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: datastore\|numberWriteAveraged_average |
| Datastore\|Read Throughput (KBps) | Rate of reading data from the datastore in kilobytes per second.<br>Key: datastore\|read_average |
| Datastore\|Storage I/O Control normalized latency (ms) | Normalized latency in microseconds on the datastore. Data for all virtual machines is combined.<br>Key: datastore\|sizeNormalizedDatastoreLatency_average |
| Datastore\|Read Latency (ms) | Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency.<br>Key: datastore\|totalReadLatency_average |
| Datastore\|Write Latency (ms) | Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency.<br>Key: datastore\|totalWriteLatency_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Datastore\|Write Throughput (KBps) | Rate of writing data to the datastore in kilobytes per second.<br>Key: datastore\|write_average |
| Datastore\|Max Queue Depth | Max Queue Depth.<br>Key: datastore\|datastoreMaxQueueDepth_latest |
| Datastore\|Highest Latency | Highest Latency.<br>Key: datastore\|maxTotalLatency_latest |
| Datastore\|Total Latency Max | Total Latency Max (ms).<br>Key: datastore\|totalLatency_max |
| Datastore\|Read Latency | Read Latency.<br>Key: datastore\|datastoreNormalReadLatency_latest |
| Datastore\|Write Latency | Write Latency.<br>Key: datastore\|datastoreNormalWriteLatency_latest |
| Datastore\|Data Read | Data Read.<br>Key: datastore\|datastoreReadBytes_latest |
| Datastore\|Data Read Rate | Data Rate.<br>Key: datastore\|datastoreReadIops_latest |
| Datastore\|Read Load | Storage DRS metric read load.<br>Key: datastore\|datastoreReadLoadMetric_latest |
| Datastore\|Outstanding Read Requests | Outstanding Read Requests.<br>Key: datastore\|datastoreReadOIO_latest |
| Datastore\|Data Written | Data Written.<br>Key: datastore\|datastoreWriteBytes_latest |
| Datastore\|Data Write Rate | Data Write Rate.<br>Key: datastore\|datastoreWriteIops_latest |
| Datastore\|Write Load | Storage DRS metric write load.<br>Key: datastore\|datastoreWriteLoadMetric_latest |
| Datastore\|Outstanding Write Requests | Outstanding Write Requests.<br>Key: datastore\|datastoreWriteOIO_latest |
| Datastore\|VM Disk I/O Workload Disparity | Percentage Disk I/O workload disparity among the VMs on the Host.<br>Key: datastore\|vmWorkloadDisparityPc |
| Datastore\|Peak Datastore Read Latency | The highest read latency among the datastores. A high number indicates that one or more datastores are experiencing poor performance.<br>Key: datastore\|peak_datastore_readLatency |
| Datastore\|Peak Datastore Write Latency | The highest write latency among the datastores. A high number indicates that one or more datastores are experiencing poor performance.<br>Key: datastore\|peak_datastore_writeLatency |

## Disk Metrics for Host Systems

Disk metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| Disk\|Total Throughput (KBps) | Average of the sum of the data read and written for all of the disk instances of the host or virtual machine.<br>disk\|usage_average |
| Disk\|I/O Usage Capacity | This metric is a function of storage\|usage_average and disk\|workload. storage\|usage_average is an average over all storage devices. This means that disk\|usage_capacity is not specific to the selected VM or the host of the VM.<br><br>Key: disk\|usage_capacity |
| Disk\|Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: disk\|commandsAveraged_average |
| Disk\|Total Latency (ms) | The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency.<br>Key: disk\|totalLatency_average |
| Disk\|Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: disk\|numberReadAveraged_average |
| Disk\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: disk\|numberWriteAveraged_average |
| Disk\|Read Throughput (KBps) | Amount of data read in the performance interval.<br>Key: disk\|read_average |
| Disk\|Write Throughput (KBps) | Amount of data written to disk in the performance interval.<br>Key: disk\|write_average |
| Disk\|Bus Resets | The number of bus resets in the performance interval.<br>Key: disk\|busResets_summation |
| Disk\|Read Latency (ms) | The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency.<br>Key: disk\|totalReadLatency_average |
| Disk\|Write Latency (ms) | The average amount of time taken for a write from the perspective of a Guest OS. This is the sum of Kernel Write Latency and Physical Device Write Latency.<br>Key: disk\|totalWriteLatency_average |
| Disk\|Physical Device Latency (ms) | The average time taken to complete a command from the physical device.<br>Key: disk\|deviceLatency_average |
| Disk\|Kernel Latency (ms) | The average time spent in ESX Server VMKernel per command.<br>Key: disk\|kernelLatency_average |
| Disk\|Queue Latency (ms) | The average time spent in the ESX Server VMKernel queue per command.<br>Key: disk\|queueLatency_average |
| Disk\|Number of Outstanding IO Operations | Number of Outstanding IO Operations.<br>Key: disk\|diskoio |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Disk\|Queued Operations | Queued Operations.<br>Key: disk\|diskqueued |
| Disk\|Demand | Demand.<br>Key: disk\|diskdemand |
| Disk\|Total Queued Outstanding operations | Sum of Queued Operation and Outstanding Operations.<br>Key: disk\|sum_queued_oio |
| Disk\|Max Observed OIO | Max Observed IO for a disk.<br>Key: disk\|max_observed |
| Disk\|Highest Latency | Highest Latency.<br>Key: disk\|maxTotalLatency_latest |
| Disk\|Max Queue Depth | Maximum queue depth during the collection interval.<br>Key: disk\|maxQueueDepth_average |
| Disk\|SCSI Reservation Conflicts | SCSI Reservation Conflicts.<br>Key: disk\|scsiReservationConflicts_summation |

**Memory Metrics for Host Systems**

Memory metrics provide information about memory use and allocation.

| Metric Name | Description |
|---|---|
| Mem\|Contention (%) | This metric is used to monitor ESXi memory usage. When the value is high, it means the ESXi is using a good percentage of available memory. You may need to add more memory to other memory-related metrics.<br><br>Key: mem\|host_contentionPct |
| Mem\|Contention (KB) | Host contention in kilobytes.<br>Key: mem\|host_contention |
| Mem\|Host Usage (KB) | Machine usage in kilobytes.<br>Key: mem\|host_usage |
| Mem\|Machine Demand (KB) | Host demand in kilobytes.<br>Key: mem\|host_demand |
| Mem\|Overall Memory used to run VMs on Host (KB) | Overall memory used to run virtual machines on the host in kilobytes.<br>Key: mem\|host_usageVM |
| Mem\|Provisioned Memory (KB) | Provisioned memory in kilobytes.<br>Key: mem\|host_provisioned |
| Mem\|Minimum Free Memory (KB) | Minimum free memory.<br>Key: mem\|host_minfree |
| Mem\|Reserved Capacity (%) | Percent reserved capacity.<br>Key: mem\|reservedCapacityPct |
| Mem\|Usable Memory (KB) | Usable memory in kilobytes.<br>Key: mem\|host_usable |
| Mem\|Usage (%) | Memory currently in use as a percentage of total available memory. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: mem\|host_usagePct |
| Mem\|ESX System Usage | Memory usage by the VMkernel and ESX user-level services.<br>Key: mem\|host_systemUsage |
| Mem\|Guest Active (KB) | Amount of memory that is actively used.<br>Key: mem\|active_average |
| Mem\|Consumed (KB) | Amount of host memory consumed by the virtual machine for guest memory.<br>Key: mem\|consumed_average |
| Mem\|Granted (KB) | Amount of memory available for use.<br>Key: mem\|granted_average |
| Mem\|Heap (KB) | Amount of memory allocated for heap.<br>Key: mem\|heap_average |
| Mem\|Heap Free (KB) | Amount of free space in the heap.<br>Key: mem\|heapfree_average |
| Mem\|VM Overhead (KB) | Memory overhead reported by host.<br>Key: mem\|overhead_average |
| Mem\|Reserved Capacity (KB) | Sum of memory reserved by consumers - Resource Pools (RP) and powered on VMs that are not in Resource Pools. RP reservations at the vCenter level are distributed to the hosts depending on the number of powered on VMs and their entitlement within the RP. For more information on resource allocation reservation, share, and limit, see Configuring Resource Allocation Settings.<br>Key: mem\|reservedCapacity_average |
| Mem\|Shared (KB) | Amount of shared memory in kilobytes.<br>Key: mem\|shared_average |
| Mem\|Shared Common (KB) | Amount of shared common memory in kilobytes.<br>Key: mem\|sharedcommon_average |
| Mem\|Swap In (KB) | Amount of memory swapped in.<br>Key: mem\|swapin_average |
| Mem\|Swap Out KB) | Amount of memory swapped out.<br>Key: mem\|swapout_average |
| Mem\|Swap Used (KB) | Amount of memory used for swapped space in kilobytes.<br>Key: mem\|swapused_average |
| Mem\|VM kernel Usage (KB) | Amount of memory used by the VM kernel.<br>Key: mem\|sysUsage_average |
| Mem\|Unreserved (KB) | Amount of unreserved memory in kilobytes.<br>Key: mem\|unreserved_average |
| Mem\|Balloon (KB) | This metric shows the total amount of memory currently used by the VM memory control. This memory was reclaimed from the respective VMs at some point in the past, and was not returned.<br>Use this metric to monitor how much VM memory has been reclaimed by ESXi through memory ballooning. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | The presence of ballooning indicates the ESXi has been under memory pressure. The ESXi activates ballooning when consumed memory reaches a certain threshold.<br><br>Look for increasing size of ballooning. This indicates that there has been a shortage of memory more than once. Look for size fluctuations which indicate the ballooned out page was actually required by the VM. This translates into a memory performance problem for the VM requesting the page, since the page must first be brought back from the disk.<br><br>Key: mem\|vmmemctl_average |
| Mem\|Zero (KB) | Amount of memory that is all zero.<br>Key: mem\|zero_average |
| Mem\|State (0-3) | Overall state of the memory. The value is an integer between 0 (high) and 3 (low).<br>Key: mem\|state_latest |
| Mem\|Usage (KB) | Host memory use in kilobytes.<br>Key: mem\|host_usage |
| Mem\|Usage (%) | Memory currently in use as a percentage of total available memory.<br>Key: mem\|usage_average |
| Mem\|Swap In Rate (KBps) | Rate at which memory is swapped from disk into active memory during the interval in kilobyte per second.<br>Key: mem\|swapinRate_average |
| Mem\|Swap Out Rate (KBps) | Rate at which memory is being swapped from active memory to disk during the current interval in kilobytes per second.<br>Key: mem\|swapoutRate_average |
| Mem\|Active Write (KB) | Average active writes in kilobytes.<br>Key: mem\|activewrite_average |
| Mem\|Compressed (KB) | Average memory compression in kilobytes.<br>Key: mem\|compressed_average |
| Mem\|Compression Rate (KBps) | Average compression rate in kilobytes per second.<br>Key: mem\|compressionRate_average |
| Mem\|Decompression Rate (KBps) | Decompression rate in kilobytes per second.<br>Key: mem\|decompressionRate_average |
| Mem\|Total Capacity (KB) | Sum of the amount of physical memory configured on ESXi hosts of the cluster in KB.<br>Key: mem\|host_provisioned |
| Mem\|Latency | Percentage of time the VM is waiting to access swapped or compressed memory.<br>Key: mem\|latency_average |
| Mem\|Capacity Contention | Capacity Contention.<br>Key: mem\|capacity.contention_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Mem\|Swap In Rate from Host Cache | Rate at which memory is being swapped from host cache into active memory.<br>Key: mem\|llSwapInRate_average |
| Mem\|Swap In from Host Cache | Amount of memory swapped-in from host cache.<br>Key: mem\|llSwapIn_average |
| Mem\|Swap Out Rate to Host Cache | Rate at which memory is being swapped to host cache from active memory.<br>Key: mem\|llSwapOutRate_average |
| Mem\|Swap Out to Host Cache | Amount of memory swapped-out to host cache.<br>Key: mem\|llSwapOut_average |
| Mem\|Swap Space Used in Host Cache | Space used for caching swapped pages in the host cache.<br>Key: mem\|llSwapUsed_average |
| Mem\|Low Free Threshold | Threshold of free host physical memory below which ESX begins to reclaim memory from VMs through ballooning and swapping.<br>Key: mem\|lowfreethreshold_average |
| Mem\|VM Memory Workload Disparity | Percentage Memory workload disparity among the VMs on the Host.<br>Key: mem\|vmWorkloadDisparityPct |
| Mem\|Active Host Load For Balance (Long Term) | Active Host Load For Balance (Long Term).<br>Key: mem\|active_longterm_load |
| Mem\|Active Host Load For Balance (Short Term) | Active Host Load For Balance (Short Term).<br>Key: mem\|active_shortterm_load |
| Mem\|Utilization | Memory utilization level based on descendant Virtual Machines utilization. Includes reservations, limits and overhead to run Virtual Machines<br>Key: mem\|total_need |

**Network Metrics for Host Systems**

Network metrics provide information about network performance.

| Metric Name | Description |
|---|---|
| Network\|Driver | This metric displays the type of network driver.<br>Key: net\|driver |
| Network\|Speed | This metric displays the network speed.<br>Key: net\|speed |
| Network\|Management Address | This metric displays the management address of the host network.<br>Key: net\|management address |
| Network\|IP Address | This metric displays the IP address of the host network.<br>Key: net\|IPaddress |
| Net\|Packets Transmitted per second | This metric shows the number of packets transmitted during the collection interval.<br>Key: net\|packetsTxPerSec |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Net\|Packets per second | Number of packets transmitted and received per second.<br>Key: net\|packetsPerSec |
| Net\|Total Throughput (KBps) | The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.<br>Key: net\|usage_average |
| Net\|I/O Usage Capacity | I/O Usage Capacity.<br>Key: net\|usage_capacity |
| Net\|Data Transmit Rate (KBps) | Average amount of data transmitted per second.<br>Key: net\|transmitted_average |
| Net\|Data Receive Rate (KBps) | Average amount of data received per second.<br>Key: net\|received_average |
| Net\|Packets Received | Number of packets received in the performance interval.<br>Key: net\|packetsRx_summation |
| Net\|Packets Transmitted | Number of packets transmitted in the performance interval.<br>Key: net\|packetsTx_summation |
| Net\|Broadcast Packets Received | Number of broadcast packets received during the sampling interval.<br>Key: net\|broadcastRx_summation |
| Net\|Broadcast Packets Transmitted | Number of broadcast packets transmitted during the sampling interval.<br>Key: net\|broadcastTx_summation |
| Net\|Error Packets Transmitted | Number of packets with errors transmitted.<br>Key: net\|errorsTx_summation |
| Net\|Multicast Packets Received | Number of multicast packets received.<br>Key: net\|multicastRx_summation |
| Net\|Multicast Packets Transmitted | Number of multicast packets transmitted.<br>Key: net\|multicastTx_summation |
| Net\|FT Throughput Usage | FT Throughput Usage.<br>Key: net\|throughput.usage.ft_average |
| Net\|HBR Throughput Usage | HBR Throughput Usage.<br>Key: net\|throughput.usage.hbr_average |
| Net\|iSCSI Throughput Usage | iSCSI Throughput Usage.<br>Key: net\|throughput.usage.iscsi_average |
| Net\|NFS Throughput Usage | NFS Throughput Usage.<br>Key: net\|throughput.usage.nfs_average |
| Net\|VM Throughput Usage | VM Throughput Usage.<br>Key: net\|throughput.usage.vm_average |
| Net\|vMotion Throughput Usage | vMotion Throughput Usage.<br>Key: net\|throughput.usage.vmotion_average |
| Net\|Unknown Protocol Frames Received | Number of frames with unknown protocol received.<br>Key: net\|unknownProtos_summation |

## System Metrics for Host Systems

System metrics provide information about the amount of CPU that resources and other applications use.

| Metric Name | Description |
|---|---|
| Sys\|Power On | 1 if the host system is powered on, 0 if the host system is powered off, or -1 if the power state is unknown.<br>Key: sys\|poweredOn |
| Sys\|Uptime (seconds) | Number of seconds since the last system startup.<br>Key: sys\|uptime_latest |
| Sys\|Disk Usage (%) | Percent disk use.<br>Key: sys\|diskUsage_latest |
| Sys\|Resource CPU Usage (MHz) | Amount of CPU that the Service Console and other applications use.<br>Key: sys\|resourceCpuUsage_average |
| Sys\|Resource CPU Active (1 min. average) | Percentage of resource CPU that is active. Average value during a one-minute period.<br>Key: sys\|resourceCpuAct1_latest |
| Sys\|Resource CPU Active (%) (5 min. average) | Percentage of resource CPU that is active. Average value during a five-minute period.<br>Key: sys\|resourceCpuAct5_latest |
| Sys\|Resource CPU Alloc Max (MHz) | Maximum resource CPU allocation in megahertz.<br>Key: sys\|resourceCpuAllocMax_latest |
| Sys\|Resource CPU Alloc Min (MHz) | Minimum resource CPU allocation in megahertz.<br>Key: sys\|resourceCpuAllocMin_latest |
| Sys\|Resource CPU Alloc Shares | Number of resource CPU allocation shares.<br>Key: sys\|resourceCpuAllocShares_latest |
| Sys\|Resource CPU Max Limited (%) (1 min. average) | Percent of resource CPU that is limited to the maximum amount. Average value during a one-minute period.<br>Key: sys\|resourceCpuMaxLimited1_latest |
| Sys\|Resource CPU Max Limited (%) (5 min. average) | Percentage of resource CPU that is limited to the maximum amount. Average value during a five-minute period.<br>Key: sys\|resourceCpuMaxLimited5_latest |
| Sys\|Resource CPU Run1 (%) | Percent resource CPU for Run1.<br>Key: sys\|resourceCpuRun1_latest |
| Sys\|Resource CPU Run5 (%) | Percent resource CPU for Run5.<br>Key: sys\|resourceCpuRun5_latest |
| Sys\|Resource Memory Alloc Max (KB) | Maximum resource memory allocation in kilobytes.<br>Key: sys\|resourceMemAllocMax_latest |
| Sys\|Resource Memory Alloc Min (KB) | Minimum resource memory allocation in kilobytes.<br>Key: sys\|resourceMemAllocMin_latest |
| Sys\|Resource Memory Alloc Shares | Number of resource memory shares allocated.<br>Key: sys\|resourceMemAllocShares_latest |
| Sys\|Resource Memory Cow (KB) | Cow resource memory in kilobytes.<br>Key: Sys\|resourceMemCow_latest |
| Sys\|Resource Memory Mapped (KB) | Mapped resource memory in kilobytes.<br>Key: ys\|resourceMemMapped_latest |
| Sys\|Resource Memory Overhead (KB) | Resource memory overhead in kilobytes.<br>Key: sys\|resourceMemOverhead_latest |
| Sys\|Resource Memory Shared (KB) | Shared resource memory in kilobytes. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: sys\|resourceMemShared_latest |
| Sys\|Resource Memory Swapped (KB) | Swapped resource memory in kilobytes.<br>Key: sys\|resourceMemSwapped_latest |
| Sys\|Resource Memory Touched (KB) | Touched resource memory in kilobytes.<br>Key: sys\|resourceMemTouched_latest |
| Sys\|Resource Memory Zero (KB) | Zero resource memory in kilobytes.<br>Key: sys\|resourceMemZero_latest |
| Sys\|Resource Memory Consumed | Resource Memory Consumed Latest (KB).<br>Key: sys\|resourceMemConsumed_latest |
| Sys\|Resource File descriptors usage | Resource File descriptors usage (KB).<br>Key: sys\|resourceFdUsage_latest |
| Sys\|vMotion Enabled | 1 if vMotion is enabled or 0 if vMotion is not enabled.<br>Key: sys\|vmotionEnabled |
| Sys\|Not in Maintenance | Not in maintenance.<br>Key: sys\|notInMaintenance |

### Management Agent Metrics for Host Systems

Management agent metrics provide information about memory use.

| Metric Name | Description |
|---|---|
| Management Agent\|Memory Used (%) | Amount of total configured memory that is available for use.<br>Key: managementAgent\|memUsed_average |
| Management Agent\|Memory Swap Used (KB) | Sum of the memory swapped by all powered-on virtual machines on the host.<br>Key: managementAgent\|swapUsed_average |
| Management Agent\|Memory Swap In (KBps) | Amount of memory that is swapped in for the Service Console.<br>Key: managementAgent\|swapIn_average |
| Management Agent\|Memory Swap Out (KBps) | Amount of memory that is swapped out for the Service Console.<br>Key: managementAgent\|swapOut_average |
| Management Agent\|CPU Usage | CPU usage.<br>Key: managementAgent\|cpuUsage_average |

### Storage Adapter Metrics for Host Systems

Storage adapter metrics provide information about data storage use.

| Metric Name | Description |
|---|---|
| Storage Adapter\|Driver | Displays the driver details of the storage adapter.<br>Key: storage adapter\|driver |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Storage Adapter\|Port WWN | Displays the world wide network port for the storage adapter.<br>Key: storage adapter\|portwwn |
| Storage Adapter\|Total Usage (KBps) | Total latency.<br>Key: storageAdapter\|usage |
| Storage Adapter\|Total IOPS | Average number of commands issued per second by the storage adapter during the collection interval.<br>Key: storageAdapter\|commandsAveraged_average |
| Storage Adapter\|Read IOPS | Average number of read commands issued per second by the storage adapter during the collection interval.<br>Key: storageAdapter\|numberReadAveraged_average |
| Storage Adapter\|Write IOPS | Average number of write commands issued per second by the storage adapter during the collection interval.<br>Key: storageAdapter\|numberWriteAveraged_average |
| Storage Adapter\|Read Throughput (KBps) | Rate of reading data by the storage adapter.<br>Key: storageAdapter\|read_average |
| Storage Adapter\|Read Latency (ms) | This metric shows the average amount of time for a read operation by the storage adapter.<br>Use this metric to monitor the storage adapter read operation performance. A high value means that the ESXi is performing a slow storage read operation.<br><br>Total latency is the sum of kernel latency and device latency.<br>Key: storageAdapter\|totalReadLatency_average |
| Storage Adapter\|Write Latency (ms) | This metric shows the average amount of time for a write operation by the storage adapter.<br>Use this metric to monitor the storage adapter write performance operation. A high value means that the ESXi is performing a slow storage write operation.<br><br>Total latency is the sum of kernel latency and device latency.<br>Key: storageAdapter\|totalWriteLatency_average |
| Storage Adapter\|Write Throughput (KBps) | Rate of writing data by the storage adapter.<br>Key: storageAdapter\|write_average |
| Storage Adapter\|Demand | Demand.<br>Key: storageAdapter\|demand |
| Storage Adapter\|Highest Latency | Highest Latency.<br>Key: torageAdapter\|maxTotalLatency_latest |
| Storage Adapter\|Outstanding Requests | Outstanding Requests.<br>Key: storageAdapter\|outstandingIOs_average |
| Storage Adapter\|Queue Depth | Queue Depth.<br>Key: storageAdapter\|queueDepth_average |
| Storage Adapter\|Queue Latency (ms) | The average time spent in the ESX Server VM Kernel queue per command.<br>Key: storageAdapter\|queueLatency_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Storage Adapter\|Queued | Queued.<br>Key: storageAdapter\|queued_average |
| Storage Adapter\|Peak Adapter Read Latency | The highest read latency among the storage adapters. A high number indicates that one or more storage adapters are experiencing poor performance.<br>Key: storageAdapter\|peak_adapter_readLatency |
| Storage Adapter\|Peak Adapter Write Latency | The highest write latency among the storage adapters. A high number indicates that one or more storage adapters are experiencing poor performance.<br>Key: storageAdapter\|peak_adapter_writeLatency |

## Storage Metrics for Host Systems

Storage metrics provide information about storage use.

| Metric Name | Description |
|---|---|
| Storage\|Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: storage\|commandsAveraged_average |
| Storage\|Read Latency (ms) | Average amount of time for a read operation in milliseconds.<br>Key: storage\|totalReadLatency_average |
| Storage\|Read Throughput (KBps) | Read throughput rate in kilobytes.<br>Key: storage\|read_average |
| Storage\|Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: storage\|numberReadAveraged_average |
| Storage\|Total Latency (ms) | Total latency in milliseconds.<br>Key: storage\|totalLatency_average |
| Storage\|Total Usage (KBps) | Total throughput rate in kilobytes per second.<br>Key: storage\|usage_average |
| Storage\|Write Latency (ms) | Average amount of time for a write operation in milliseconds.<br>Key: storage\|totalWriteLatency_average |
| Storage\|Write Throughput (KBps) | Write throughput rate in kilobytes per second.<br>Key: storage\|write_average |
| Storage\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: storage\|numberWriteAveraged_average |

## Sensor Metrics for Host Systems

Sensor metrics provide information about host system cooling.

| Metric Name | Description |
|---|---|
| Sensor\|Fan\|Speed (%) | Percent fan speed.<br>Key: Sensor\|fan\|currentValue |
| Sensor\|Fan\|Health State | Fan health state.<br>Key: Sensor\|fan\|healthState |
| Sensor\|Temperature\|Temp C | Fan temperature in centigrade.<br>Key: Sensor\|temperature\|currentValue |
| Sensor\|Temperature\|Health State | Fan health state.<br>Key: Sensor\|temperature\|healthState |

**Power Metrics for Host Systems**

Power metrics provide information about host system power use.

| Metric Name | Description |
|---|---|
| Power\|Total Energy Consumed in the collection period (Wh) | Displays the total electricity consumed based on the time interval selected. The default collection cycle is set to 5 mins. You can continue using the default setting or edit it for each adapter instance. For example, if the time interval is set to its default value, the value represents the energy consumed per 5 mins. |
| Power\|Total Host System Power Consumed in an Hour (Wh) | Displays the the total electricity power consumed in an hour by ESXi Host. The data collected is over a period of an hour and published along with the other metrics in VMware Aria OperationsVMware Cloud Foundation Operations. In case of a connectivity or availability issue in VMware Aria OperationsVMware Cloud Foundation Operations or vCenter adapter instance, this hourly metric might not be published and the missed value during this period does not get recalculated. Once the connection is re-established, the next data points get published.<br>**NOTE**<br>This metric is deactivated by default. You can activate it from the Policies page. For more information, see Metrics and Properties Details in the *VMware Aria OperationsVMware Cloud Foundation Operations Configuration Guide*.This metric is deactivated by default. You can activate it from the Policies page. For more information, see Metrics and Properties Details. |
| Power\|Power (Watt) | Host power use in watts.<br>Key: power\|power_average |
| Power\|Current Power Consumption Rate (Watt) | The power consumption rate per second, averaged over the reporting period.<br>Key: power\|power_average |
| Power\|Power Cap (Watt) | Host power capacity in watts. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: power\|powerCap_average |
| Power\|Host Power Capacity Usage – Idle | Power consumed by the host in its idle state. This is the power consumed by the host when there are no VMs in it.<br>Key: power\|capacity.usageIdle_average |
| Power\|(DEP) Energy (Joule) | Total energy consumed in joules.<br>Key: power\|energy_summation |

**Disk Space Metrics for Host Systems**

Disk space metrics provide information about disk space use.

| Metric Name | Description |
|---|---|
| Diskspace\|Number of Virtual Disks | Number of virtual disks.<br>Key: diskspace\|numvmdisk |
| Diskspace\|Shared Used (GB) | Used shared disk space in gigabytes.<br>Key: diskspace\|shared |
| Diskspace\|Snapshot | Disk space used by snapshots in gigabytes.<br>Key: diskspace\|snapshot |
| Diskspace\|Virtual Disk Used (GB) | Disk space used by virtual disks in gigabytes.<br>Key: diskspace\|diskused |
| Diskspace\|Virtual machine used (GB) | Disk space used by virtual machines in gigabytes.<br>Key: diskspace\|used |
| Diskspace\|tTotal disk space used | Total disk space used on all datastores visible to this object.<br>Key: diskspace\|total_usage |
| Diskspace\|Total disk spacey | Total disk space on all datastores visible to this object.<br>Key: diskspace\|total_capacity |
| Diskspace\|Total provisioned disk space | Total provisioned disk space on all datastores visible to this object.<br>Key: diskspace\|total_provisioned . |
| Diskspace\|Utilization (GB) | Storage space utilized on connected vSphere datastores.<br>Key: diskspace\|total_usage |
| Diskspace\|Workload (%) | Total storage space available on connected vSphere datastores.<br>Key: diskspace\|total_capacity |

**Summary Metrics for Host Systems**

Summary metrics provide information about overall host system performance.

| Metric Name | Description |
|---|---|
| Summary\|Number of Running VMs | This metric shows the number of VMs running on the host during the last metric collection time.<br>Large spikes of running VMs might be a reason for CPU or memory spikes as more resources are used in the host.<br><br>Number of Running VMs gives you a good indicator of how many requests the ESXi host must juggle. This excludes powered off VMs as they do not impact ESXi performance. A change in this number in your environment can contribute to performance problems. A high number of running VMs in a host also means a higher concentration risk, as all the VMs will become unavailable (or be relocated by HA) if the ESXi crashes.<br><br>Look for any correlation between spikes in the number of running VMs and spikes in other metrics such as CPU Contention/Memory Contention.<br><br>Key: summary\|number_running_vms |
| Summary\|Maximum Number of VMs | Maximum number of virtual machines<br>Key: summary\|max_number_vms |
| Summary\|Number of vMotions | This metric shows the number of vMotions that occurred in the host in the last X minutes.<br>The number of vMotions is a good indicator of stability. In a healthy environment, this number should be stable and relatively low.<br><br>Look for correlation between vMotions and spikes in other metrics such as CPU/Memory contention.<br><br>The vMotion should not create any spikes, however, the VMs moved into the host might create spikes in memory usage, contention and CPU demand and contention.<br><br>Key: summary\|number_vmotion |
| Summary\|Total Number of Datastores | Total Number of Datastores.<br>Key: summary\|total_number_datastores |
| Summary\|Number of VCPUs on Powered On VMs | Total number of VCPUs of Virtual Machines that are powered on.<br>Key: summary\|number_running_vcpus |
| Summary\|Total Number of VMs | Total number of virtual machines.<br><br>**NOTE**<br>This is the total number of VMs excluding VM templates.<br><br>Key: summary\|total_number_vms |
| Summary\|Number of VM Templates | Number of VM Templates<br>Key: summary\|number_vm_templates |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Summary\|Consider for Balance | Summary\|Consider for Balance = 1 when the host is Powered On, Connected, not in Maintenance Mode, and not a Failover Host, otherwise it = -1 |

**HBR Metrics for Host Systems**

Host-based replication (HBR) metrics provide information about vSphere replication.

| Metric Name | Description |
|---|---|
| HBR\|Replication Data Received Rate | Replication Data Received Rate.<br>Key: hbr\|hbrNetRx_average |
| HBR\|Replication Data Transmitted Rate | Replication Data Transmitted Rate.<br>Key: hbr\|hbrNetTx_average |
| HBR\|Replicated VM Count | Number of replicated virtual machines.<br>Key: hbr\|hbrNumVms_average |

**Cost Metrics for Host Systems**

Cost metrics provide information about the cost.

| Metric Name | Description |
|---|---|
| Monthly Maintenance Total Cost | Monthly total cost for maintenance.<br>Key: cost\|maintenanceTotalCost |
| Monthly Host OS License Total Cost | Monthly total cost for the host operating system license.<br>Key: cost\|hostOslTotalCost |
| Monthly Network Total Cost | Monthly total cost for network including cost of NIC cards associated with host.<br>Key: cost\|networkTotalCost |
| Monthly Server Hardware Total Cost | Monthly total cost for server hardware, based on amortized monthly value.<br>Key: cost\|hardwareTotalCost |
| Monthly Facilities Total Cost | Monthly total cost of facilities including real estate, power, and cooling.<br>Key: cost\|facilitiesTotalCost |
| Monthly Server Labor Total Cost | Monthly total cost for the server operating system labor.<br>Key: cost\|hostLaborTotalCost |
| Monthly Server Fully Loaded Cost | Monthly cost for a fully loaded server incorporating all cost driver values attributed to the server.<br>Key: cost\|totalLoadedCost |
| MTD Server Total Cost | Month to date cost for a fully loaded server incorporating all cost driver values attributed to the server.<br>Key: totalMTDCost |
| Server Accumulated Depreciation | Month to date accumulated cost for a deprecated server.<br>Key: Cost\|Server Accumulated Depreciation |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Aggregated Daily Total Cost | Daily aggregate daily total cost of the deleted VM present in the host system.<br>Key: Cost\|aggregatedDailyTotalCost |
| Aggregated Deleted VM Daily Total Cost | Daily aggregate cost of the deleted VM present in the host system.<br>Key: Cost\|aggregatedDeletedVmDailyTotalCost |

**Disabled Instanced Metrics**

The instance metrics created for the following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that these metrics collect data by default but all the instanced metrics created for these metrics, do not collect data by default.

| Metric Name |
|---|
| Datastore\|Outstanding IO requests (OIOs) |
| Datastore\|Read IOPS |
| Datastore\|Read Latency (ms) |
| Datastore\|Read Throughput (KBps) |
| Datastore\|Total Latency (ms) |
| Datastore\|Total Throughput (KBps) |
| Datastore\|unmapIOs_summation |
| Datastore\|unmapsize_summation |
| Datastore\|Write IOPS |
| Datastore\|Write Latency (ms) |
| Datastore\|Write Throughput (KBps) |
| Disk\|Physical Device Latency (ms) |
| Disk\|Queue Latency (ms) |
| Disk\|Read IOPS |
| Disk\|Read Latency (ms) |
| Disk\|Read Throughput (KBps) |
| Disk\|Write IOPS |
| Disk\|Write Latency (ms) |
| Disk\|Write Throughput (KBps) |
| Net\|Data Receive Rate (KBps) |
| Net\|Data Transmit Rate (KBps) |
| Net\|Error Packets Transmitted |
| Net\|Packets Dropped (%) |
| Net\|Packets Transmitted per second |
| Net\|Received Packets Dropped |
| Net\|Transmitted Packets Dropped |
| Net\|Usage Rate (%) |
| Storage Adapter\|Read IOPS |
| Storage Adapter\|Read Latency (ms) |

*Table continued on next page*

*Continued from previous page*

| Metric Name |
|---|
| Storage Adapter\|Read Throughput (KBps) |
| Storage Adapter\|Write IOPS |
| Storage Adapter\|Write Latency (ms) |
| Storage Adapter\|Write Throughput (KBps) |

**Disabled Metrics**

The following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

You can enable these metrics in the Policy workspace.

| Metric Name | Key |
|---|---|
| CPU\|Idle (msec) | cpu\|idle_summation |
| CPU\|Used (msec) | cpu\|used_summation |
| Datastore I/O\|Average Observed Virtual Machine Disk I/O Workload | datastore\|vmPopulationAvgWorkload |
| Datastore I/O\|Max Observed Number of Outstanding IO Operations | datastore\|maxObserved_OIO |
| Datastore I/O\|Max Observed Read Rate (kbps) | datastore\|maxObserved_Read |
| Datastore I/O\|Max Observed Reads per second | datastore\|maxObserved_NumberRead |
| Datastore I/O\|Max Observed Write Rate (kbps) | datastore\|maxObserved_Write |
| Datastore I/O\|Max Observed Writes per second | datastore\|maxObserved_NumberWrite |
| Datastore I/O\|Maximum Observed VM Disk I/O Workload | datastore\|vmPopulationMaxWorkload |
| Network I/O\|bytesRx (kbps) | net\|bytesRx_average |
| Network I/O\|bytesTx (kbps) | net\|bytesTx_average |
| Network I/O\|Demand (%) | net\|demand |
| Network I/O\|Error Packets Received | net\|errorsRx_summation |
| Network I/O\|Max Observed Received Throughput (kbps) | net\|maxObserved_Rx_KBps |
| Network I/O\|Max Observed Throughput (kbps) | net\|maxObserved_KBps |
| Network I/O\|Max Observed Transmitted Throughput (kbps) | net\|maxObserved_Tx_KBps |
| Network I/O\|Packets Received per second | net\|packetsRxPerSec |
| Network I/O\|Packets Dropped | net\|dropped |
| Summary\|Workload Indicator | summary\|workload_indicator |
| vFlash Module\|Latest Number of Active Vm Disks | vflashModule\|numActiveVMDKs_latest |
| Net\|Received Packets Dropped | Number of received packets dropped in the performance interval.<br>Key: net\|droppedRx_summation |
| Net\|Transmitted Packets Dropped | Number of transmitted packets dropped in the performance interval.<br>Key: net\|droppedTx_summation |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Key |
|---|---|
| Net\|Packets Dropped (%) | This metric shows the percentage of received and transmitted packets dropped during the collection interval.<br>This metric is used to monitor reliability and performance of the ESXi network. When a high value is displayed, the network is not reliable and performance suffers.<br><br>Key: net\|droppedPct |
| Diskspace\|Not Shared (GB) | Unshared disk space in gigabytes.<br>Key: diskspace\|notshared |

### Cluster Compute Resource Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, disk space, CPU use, disk, memory, network, power, and summary metrics for cluster compute resources.

Cluster Compute Resource metrics include capacity and badge metrics. See definitions in:
- Capacity Analytics Generated Metrics
- Badge Metrics

### GPU Metrics

GPU metrics provide information about the GPU usage and performance.

| Metric Name | Description |
|---|---|
| GPU\|Compute Utilization (%) | The compute utilization percentage of a GPU. |
| GPU\|Memory Usage (%) | Memory currently in use as a percentage of total available memory. |
| GPU\|Memory Used (KB) | The amount of GPU memory used in kilobytes. |
| GPU\|Number of GPUs | Number of GPUs. |
| GPU\|Total Memory (KB) | Total memory in kilobytes. |

### Cluster Metrics for ROI Dashboard

Cluster metrics provide information about the metrics in ROI dashboard.

| Metric Name | Description |
|---|---|
| Total Number Of Reclaimable Hosts | This metric displays the total number of reclaimable hosts across all vCenters.<br>Key: metric=cost\|reclaimableHostCost |
| Total Reclaimable Host Cost | This metric displays the reclaimable host cost based on the recommended size.<br>Key: cost\|reclaimableHostCost |

### Configuration Metrics for Cluster Compute Resources

Configuration metrics provide information about configuration settings.

| Metric Name | Description |
|---|---|
| Configuration\|DAS Configuration\|Admission Control Enabled | DAS configuration admission control enabled.<br>Key: configuration\|dasconfig\|AdministrationControlEnabled |
| Configuration\|DAS Configuration\|Active Admission Control Policy | DAS configuration active admission control policy.<br>Key: configuration\|dasconfig\|activeAdministrationControlPolicy |
| Configuration\|DRS Configuration\|Affinity Rules | Affinity rules for DRS configuration.<br>Key: configuration\|DRSconfiguration\|affinity rules |
| Configuration\|DRS Configuration\|Tolerance Imbalance Threshold | Displays the tolerance imbalance threshold for DRS configuration.<br>Key: configuration\|DRSconfiguration\|<br>ToleranceimbalanceThreshold |
| Configuration\|DRS Configuration\|Default DRS behavior | Displays the default DRS configuration behavior.<br>Key: configuration\|DRSconfiguration\|DefaultDRSbehaviour |
| Configuration\|DRS Configuration\|Idle Consumed Memory | Displays the idle memory consumed by DRS configuration.<br>Key: configuration\|DRSconfiguration\|IdleConsumedMemory |
| Configuration\|DRS Configuration\| DRS vMotion Rate | Displays the vMotion rate for the DRS configuration.<br>Key: configuration\|DRSconfiguration\|DRSvMotion Rate |
| Configuration\|DPM Configuration\|Default DPM behavior | Displays the default behavior for the DPM configuration.<br>Key: configuration\|DPMconfiguration\|DefaultDPMbehaviour |
| Configuration\|DPM Configuration\|DPM Enabled | Displays whehter the DPM Configuration is enabled or not.<br>Key: configuration\|DPMConfiguration\|DPMEnabled |
| Configuration\|Failover Level | DAS configuration failover level.<br>Key: configuration\|dasconfig\|failoverLevel |
| Configuration\|Active Admission Control Policy | DAS configuration active admission control policy.<br>Key: configuration\|dasconfig\|activeAdministrationControlPolicy |
| Configuration\|CPU Failover Resources Percent | Percent CPU failover resources for DAS configuration admission control policy.<br>Key: configuration\|dasconfig\|adminissionControlPolicy\|<br>cpuFailoverResourcesPercent |
| Configuration\|Memory Failover Resources Percent | Percent memory failover resources for DAS configuration admission control policy.<br>Key: configuration\|dasconfig\|adminissionControlPolicy\|<br>memoryFailoverResourcesPercent |

### Disk Space Metrics for Cluster Compute Resources

Disk space metrics provide information about disk space use.

| Metric Name | Description |
|---|---|
| DiskSpace\|Snapshot Space | Displays the disk space used by the snapshot.<br>Key: DiskSpace\|snapshot space |
| Diskspace\|Virtual machine used (GB) | Space used by virtual machine files in gigabytes.<br>Key: diskspace\|used |
| Diskspace\|Total disk space used | Total disk space used on all datastores visible to this object.<br>Key: diskspace\|total_usage |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Diskspace\|Total disk space | Total disk space on all datastores visible to this object.<br>Key: diskspace\|total_capacity |
| Diskspace\|Total provisioned disk space | Total provisioned disk space on all datastores visible to this object.<br>Key: diskspace\|total_provisioned |
| Diskspace\|Virtual Disk Used (GB) | Space used by virtual disks in gigabytes.<br>Key: diskspace\|diskused |
| Diskspace\|Snapshot Space (GB) | Space used by snapshots in gigabytes.<br>Key: diskspace\|snapshot |
| Diskspace\|Shared Used (GB) | Shared used space in gigabytes.<br>Key: diskspace\|shared |
| Diskspace\|Utilization (GB) | Storage space used on the connected vSphere Datastores.<br>Key: diskspace\|total_usage |
| Diskspace\|Total Capacity (GB) | Total storage space available on the connected vSphere datastores.<br>Key: diskspace\|total_capacity |

**CPU Usage Metrics for Cluster Compute Resources**

CPU usage metrics provide information about CPU use.

| Metric Name | Description |
|---|---|
| CPU\|Allocation\|Usable Capacity after HA and Buffer (vCPUs) | This metric shows the total capacity taking into consideration the over-commit ratio and after subtracting the CPU resources needed for HA and reserved buffer.<br>Key: cpu\|alloc\|usableCapacity |
| CPU\|Capacity Usage | This metric shows the percentage of the capacity used.<br>Key: cpu\|capacity_usagepct_average |
| CPU\|CPU Contention (%) | This metric is an indicator of the overall contention for CPU resources that occurs across the workloads in the cluster. When contention occurs, it means that some of the virtual machines are not immediately getting the CPU resources they are requesting. Use this metric to identify when a lack of CPU resources might be causing performance issues in the cluster.<br><br>This metric is the sum of the CPU contention across all hosts in the cluster averaged over two times the number of physical CPUs in the cluster to account for hyper-threading. CPU contention takes into account:<br>• CPU Ready<br>• CPU Co-stop<br>• Power management<br>• Hyper threading<br>This metric is more accurate than CPU Ready since it takes into account CPU Co-stop and Hyper threading. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | When using this metric, the number should be lower than the performance you expect. If you expect performance at 10%, then the number should be lower than 10%.<br><br>Since this value is averaged across all hosts in the cluster, you might find that some hosts have a higher CPU contention while others are lower. To ensure that vSphere spreads out the running workloads across hosts, consider enabling a fully automated DRS in the cluster.<br><br>Key: cpu\|capacity_contentionPct |
| CPU\|Demand\|Usable Capacity after HA and Buffer (MHz) | This metric shows the total capacity after subtracting the CPU resources needed for HA and reserved buffer.<br>Key: cpu\|demand\|usableCapacity |
| CPU\|Demand (%) | This metric is an indicator of the overall demand for CPU resources by the workloads in the cluster.<br>It shows the percentage of CPU resources that all the virtual machines might use if there were no CPU contention or CPU limits set. It represents the average active CPU load in the past five minutes.<br><br>Key: cpu\|demandPct |
| CPU\|Demand (MHz) | Sum of CPU utilization of all virtual machines on this cluster, including limits and VM overhead.<br>Key: cpu\|demandmhz |
| CPU\|Number of CPU Sockets | Number of CPU sockets.<br>Key: cpu\|numpackages |
| CPU\|Overall CPU Contention | Overall CPU contention in milliseconds.<br>Key: cpu\|capacity_contention |
| CPU\|Host Provisioned Capacity | Provisioned CPU capacity in megahertz.<br>Key: cpu\|capacity_provisioned |
| CPU\|Provisioned CPUs | Number of Physical CPUs (Cores).<br>Key: cpu\|corecount_provisioned |
| CPU\|Usage (MHz) | Average CPU use in megahertz.<br>Key: cpu\|usagemhz_average |
| CPU\| VM CPU Usage (Mhz) | Sum of CPU usages of all VMs in the cluster. The CPU\|Usage (Mhz) metric value of each VM is taken for summation.<br>Key: cpu\|vm_usagemhz_average |
| CPU\|Demand | CPU Demand.<br>Key: cpu\|demand_average |
| CPU\|Overhead | Amount of CPU overhead.<br>Key: cpu\|overhead_average |
| CPU\|Demand without overhead | Value of demand excluding any overhead.<br>Key: cpu\|demand_without_overhead |
| CPU\|Provisioned Capacity | Provisioned Capacity (MHz).<br>Key: cpu\|vm_capacity_provisioned |
| CPU\|Number of hosts stressed | Number of hosts stressed. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: cpu\|num_hosts_stressed |
| CPU\|Stress Balance Factor | Stress Balance Factor.<br>Key: cpu\|stress_balance_factor |
| CPU\|Lowest Provider Capacity Remaining | Lowest Provider Capacity Remaining.<br>Key: cpu\|min_host_capacity_remaining |
| CPU\|Workload Balance Factor | Workload Balance Factor.<br>Key: cpu\|workload_balance_factor |
| CPU\|Highest Provider Workload | Highest Provider Workload.<br>Key: cpu\|max_host_workload |
| CPU\|Host workload Max-Min Disparity | Difference of Max and Min host workload in the container.<br>Key: cpu\|host_workload_disparity |
| CPU\|Host stress Max-Min Disparity | Difference of Max and Min host stress in the container.<br>Key: cpu\|host_stress_disparity |
| CPU\|Total Capacity (MHz) | Total CPU resources configured on the descendant ESXi hosts.<br>Key: cpu\|capacity_provisioned |
| CPU\|Usable Capacity (MHz) | The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services.<br>Key: cpu\|haTotalCapacity_average |

**Disk Metrics for Cluster Compute Resources**

Disk metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| Disk\|Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: disk\|commandsAveraged_average |
| Disk\|Total Latency (ms) | Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Command Latency and Physical Device Command Latency metrics.<br>Key: disk\|totalLatency_average |
| Disk\|Read Latency (ms) | Average amount of time for a read operation from the virtual disk. The total latency is the sum of Kernel latency and device latency.<br>Key: disk\|totalReadLatency_average |
| Disk\|Write Latency (ms) | The average amount of time taken for a read from the perspective of a Guest OS. This is the sum of Kernel Read Latency and Physical Device Read Latency.<br>Key: disk\|totalWriteLatency_averag |
| Disk\|Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: disk\|numberReadAveraged_averag |
| Disk\|Total Throughput (KBps) | Average of the sum of the data read and written for all the disk instances of the host or virtual machine. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: disk\|usage_average |
| Disk\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: disk\|numberWriteAveraged_average |
| Disk\|Read Requests | Amount of data read from the disk during the collection interval.<br>Key: disk\|read_average |
| Disk\|Write Requests | Amount of data written to the disk during the collection interval.<br>Key: disk\|write_average |
| Disk\|Total Queued Outstanding operations | Sum of queued operation and outstanding operations.<br>Key: disk\|sum_queued_oio |
| Disk\|Max Observed OIO | Max observed outstanding IO for a disk.<br>Key: disk\|max_observed |

**Memory Metrics for Cluster Compute Resources**

Memory metrics provide information about memory use and allocation.

| Metric Name | Description |
|---|---|
| Mem\|Active Write (KB) | Active writes in kilobytes.<br>Key: mem\|activewrite_average |
| Mem\|Compressed (KB) | Average compression in kilobytes.<br>Key: mem\|compressed_average |
| Mem\|Compression Rate (KBps) | Average compression rate in kilobytes.<br>Key: mem\|compressionRate_average |
| Mem\|Consumed (KB) | Amount of host memory consumed by the virtual machine for guest memory.<br>Key: mem\|consumed_average |
| Mem\|Contention (%) | This metric is an indicator of the overall contention for memory resources that occurs across the workloads in the cluster. When contention occurs, it means that some of the VMs are not immediately getting the memory resources that they are requesting.<br>Use this metric to identify when lack of memory resources might be causing performance issues in the cluster.<br><br>Key: mem\|host_contentionPct |
| Mem\|Contention (KB) | Contention in kilobytes.<br>Key: mem\|host_contention |
| Mem\|Decompression Rate (KBps | Decompression rate in kilobytes.<br>Key: mem\|decompressionRate_average |
| Mem\|Granted (KB) | Amount of memory available for use.<br>Key: mem\|granted_average |
| Mem\|Guest Active (KB) | Amount of memory that is actively used.<br>Key: mem\|active_average |
| Mem\|Heap (KB) | Amount of memory allocated for heap. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: mem\|heap_average |
| Mem\|Heap Free (KB) | Free space in the heap.<br>Key: mem\|heapfree_average |
| Mem\|Balloon | This metric shows the amount of memory currently used by the virtual machine memory control. It is only defined at the VM level.<br>Key: mem\|vmmemctl_average |
| Mem\|VM Overhead (KB) | Memory overhead reported by host.<br>Key: mem\|overhead_average |
| Mem\|Provisioned Memory (KB) | Provisioned memory in kilobytes.<br>Key: mem\|host_provisioned |
| Mem\|Reserved Capacity (KB) | Sum of memory reservations by consumers such as Resource Pools (RP) and powered on VMs that are not in Resource Pools, aggregated and published on cluster level. RP reservations at the vCenter level are distributed to hosts depending on the number of powered on VMs and their entitlement within the Resource Pool. For more information on resource allocation reservation, share, and limit, see Configuring Resource Allocation Settings.<br>Key: mem\|reservedCapacity_average |
| Mem\|Shared (KB) | Amount of shared memory.<br>Key: mem\|shared_average |
| Mem\|Shared Common (KB) | Amount of shared common memory.<br>Key: mem\|sharedcommon_average |
| Mem\|Swap In (KB) | Amount of memory that is swapped in for the service console.<br>Key: mem\|swapin_average |
| Mem\|Swap In Rate (KBps) | Rate at which memory is swapped from disk into active memory during the interval.<br>Key: mem\|swapinRate_average |
| Mem\|Swap Out (KB) | Amount of memory that is swapped out for the service console.<br>Key: mem\|swapout_average |
| Mem\|Swap Out Rate (KBps) | Rate at which memory is being swapped from active memory into disk during the current interval.<br>Key: mem\|swapoutRate_average |
| Mem\|Swap Used (KB) | Amount of memory used for swap space.<br>Key: mem\|swapused_average |
| Mem\|Total Capacity (KB) | Total capacity in kilobytes.<br>Key: mem\|totalCapacity_average |
| Mem\|Unreserved (KB) | Memory available for reservation by VMs and Resource Pools (RP), aggregated and published on cluster level. The reservations are set for VMs and RPs and not for ESXi hosts or clusters. For more information on resource allocation reservation, share, and limit, see Configuring Resource Allocation Settings.<br>Key: mem\|unreserved_average |
| Mem\|Usable Memory (KB) | Usable memory in kilobytes. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: mem\|host_usable |
| Mem\|Usage/Usable | Percent memory used.<br>Key: mem\|host_usagePct |
| Mem\|Host Usage (KB) | Memory use in kilobytes.<br>Key: mem\|host_usage |
| Mem\|Machine Demand | Memory Machine Demand in KB.<br>Key: mem\|host_demand |
| Mem\|ESX System Usage | Memory usage by the VMkernel and ESX user-level services.<br>Key: mem\|host_systemUsage |
| Mem\|Usage (%) | This metric shows the portion of the total memory in all hosts in the cluster that is being used.<br>This metric is the sum of memory consumed across all hosts in the cluster divided by the sum of physical memory across all hosts in the cluster.<br>∑ memory consumed on all hosts<br>- X 100%<br>∑ physical memory on all hosts |
| Memory \|Workload (%) | Demand over usable capacity. Wherever applicable, demand includes limit and contention. |
| Mem\|Usage (KB) | Memory currently in use as a percentage of total available memory.<br>Key: mem\|usage_average |
| Mem\|VM kernel Usage (KB) | Amount of memory that the VM kernel uses.<br>Key: mem\|sysUsage_average |
| Mem\|Zero (KB) | Amount of memory that is all 0.<br>Key: mem\|zero_average |
| Mem\|Number of Hosts Stressed | Number of hosts stressed.<br>Key: mem\|num_hosts_stressed |
| Mem\|Stress Balance Factor | Stress balance factor.<br>Key: mem\|stress_balance_factor |
| Mem\|Lowest Provider Capacity Remaining | Lowest provider capacity remaining.<br>Key: mem\|min_host_capacity_remaining |
| Mem\|Workload Balance Factor | Workload balance factor.<br>Key: mem\|workload_balance_factor |
| Mem\|Highest Provider Workload | Highest provider workload.<br>Key: mem\|max_host_workload |
| Mem\|Host workload Max-Min Disparity | Difference of Max and Min host workload in the container.<br>Key: mem\|host_workload_disparity |
| Mem\|Host stress Max-Min Disparity | Difference of Max and Min host stress in the container.<br>Key: mem\|host_stress_disparity |
| Mem\|Utilization (KB) | Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: mem\|total_need |
| Mem\|Total Capacity (KB) | Sum of the amount of physical memory configured on ESXi hosts of the cluster.<br>Key: mem\|host_provisioned |
| Mem\|Usable Capacity (KB) | Memory resources available after subtracting reservations for vSphere HA (failover host/s) from the cluster's total memory capacity.<br>Hosts that are in the maintenance mode are not included in the calculation.<br><br>Key: mem\|haTotalCapacity_average |

**Network Metrics for Cluster Compute Resources**

Network metrics provide information about network performance.

| Metric Name | Description |
|---|---|
| Net\|Data Receive Rate (KBps) | Average amount of data received per second.<br>Key: net\|received_average |
| Net\|Data Transmit Rate (KBps) | Average amount of data transmitted per second.<br>Key: net\|transmitted_average |
| Net\|Packets Dropped | Number of packets dropped in the performance interval.<br>Key: net\|dropped |
| Net\|Packets Dropped (%) | Percentage of packets dropped.<br>Key: net\|droppedPct |
| Net\|Packets Received | Number of packets received in the performance interval.<br>Key: net\|packetsRx_summation |
| Net\|Packets Transmitted | Number of packets transmitted in the performance interval.<br>Key: net\|packetsTx_summation |
| Net\|Received Packets Dropped | Number of received packets dropped in the performance interval.<br>Key: net\|droppedRx_summation |
| Net\|Transmitted Packets Dropped | Number of transmitted packets dropped in the performance interval.<br>Key: net\|droppedTx_summation |
| Net\|Total Throughput (KBps) | The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.<br>Key: net\|usage_average |
| Network\|Error Packets | Displays the total number of error packets (transmitted and received) from all the ESXi in the cluster within a time interval of 20 seconds.<br>Key: Network\|Error Packets |

**Datastore Metrics for Cluster Compute Resources**

Datastore metrics provide information about Datastore use.

| Metric Name | Description |
|---|---|
| Datastore\|TotalThroughput | Displays the total throughput for the datastore.<br>Key: datastore\|thorughput |
| Datastore\|Outstanding IO requests | OIO for datastore.<br>Key: datastore\|demand_oio |
| Datastore\|Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: datastore\|numberReadAveraged_average |
| Datastore\|Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: datastore\|numberWriteAveraged_average |
| Datastore\|Read Throughput (KBps) | Amount of data read in the performance interval.<br>Key: datastore\|read_average |
| Datastore\|Write Throughput (KBps) | Amount of data written to disk in the performance interval.<br>Key: datastore\|write_average |
| Datastore\|Read Latency | Average amount of time taken for a read operation from the datastore.<br>Key: datastore\|ReadLatency |
| Datastore\|Write Latency | Average amount of time taken for a write operation from the datastore.<br>Key: datastore\|WriteLatency |
| Datastore\|Max VM Disk Latency | Maximum amount of time taken to read or write data from a virtual machine.<br>Key: datastore\|MaxVMDiskLatency |
| Datastore\|Outstanding IO Requests (OIOs) | This metric displays the outstanding datastore IO requests.<br>Key: datastore\|OutstandingIORequests |
| Datastore\|Host SCSI Disk Partition | This metric displays the datastore host scsi partition.<br>Key: datastore\|HostSCSIDiskPartition |
| Devices\|Command Aborted | The metric lists the stopped device commands.<br>Key: devices\|CommandAborted |

**Cluster Services Metrics for Cluster Compute Resources**

Cluster Services metrics provide information about cluster services.

| Metric Name | Description |
|---|---|
| Cluster Services\|Total Imbalance | Total imbalance in cluster services<br>Key: clusterServices\|total_imbalance |
| ClusterServices\|Effective CPU Resources (MHz) | VMware DRS effective CPU resources available.<br>Key: clusterServices\|effectivecpu_average |
| ClusterServices\|Effective Memory Resources (KB) | VMware DRS effective memory resources available.<br>Key: clusterServices\|effectivemem_average |
| Cluster Services\|DRS Initiated vMotion Count | clusterServices\|number_drs_vmotion |

**Power Metrics for Cluster Compute Resources**

Power metrics provide information about power use.

| Metric Name | Description |
|---|---|
| Power\|Total Energy Consumed in the collection period (Wh) | Displays the total electricity consumed based on the time interval selected. The default collection cycle is set to 5 mins. You can continue using the default setting or edit it for each adapter instance. For example, if the time interval is set to its default value, the value represents the energy consumed per 5 mins. |
| Power\|Current Power Consumption Rate (Watt) | The power consumption rate per second, averaged over the reporting period.<br>Key: power\|power_average |
| Power\|Power Cap (Watt) | Average power capacity in watts.<br>Key: power\|powerCap_average |
| Power\|(DEP) Energy (Joule) | Total energy consumed in joules<br>Key: power\|energy_summation |

**Summary Metrics for Cluster Compute Resources**

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| Summary\|Number of Running Hosts | Number of running hosts.<br>Key: summary\|number_running_hosts |
| Summary\|Number of Running VMs | This metric shows the total number of VMs running on all hosts in the cluster.<br>Key: summary\|number_running_vms |
| Summary\|Number of vMotions | This metric shows the number of vMotions that occurred during the last collection cycle.<br>When using this metric, look for a low number which indicates that the cluster might serve its VMs. A vMotion can impact VM performance during the stun time.<br><br>Key: summary\|number_vmotion |
| Summary\|Number of Hosts | Total number of hosts.<br>Key: summary\|total_number_hosts |
| Summary\|Total Number of VMs | Total number of virtual machines.<br><br>**NOTE**<br>This shows the total number of VMs excluding VM templates under the datastore.<br><br>Key: summary\|total_number_vms |
| Summary\|Total Number of Datastores | Total number of datastores.<br>Key: summary\|total_number_datastores |
| Summary\|Number of VCPUs on Powered On VMs | Number of virtual CPUs on powered-on virtual machines.<br>Key: summary\|number_running_vcpus |
| Summary\|Average Running VM Count per Running Host | Average number of running virtual machines per running host.<br>Key: summary\|avg_vm_density |
| Summary\|Cluster Availability (%) | Percentage of hosts powered-on in the cluster.<br>Key: summary\|cluster_availability |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Summary\|Datastore | Displays the status of the datastore.<br>Key: summary\|datastore |
| Summary\|Type | Displays the datastore type.<br>Key: summary\|type |
| Summary\|Is Local | Displays whether the datastore is local or not.<br>Key: summary\|Islocal |
| Summary\|Number of VM Templates | Number of VM templates.<br>Key: summary\|number_vm_templates |
| Summary\|Number of Pods | Number of pods.<br>**NOTE**<br>This is published if the cluster is Workload Management enabled or there are pods under the cluster.<br>Key: summary\|total_number_pods |
| Summary\|Number of Namespaces | Number of namespaces.<br>**NOTE**<br>This is published if the cluster is Workload Management enabled or there are namespaces under the cluster.<br>Key: summary\|numberNamespaces |
| Summary\|Number Kubernetes Clusters | Number of Kuberntes clusters.<br>**NOTE**<br>This is published if the cluster is Workload Management enabled or there are Kuberntes clusters under the cluster.<br>Key: summary\|numberKubernetesClusters |
| Summary\|Number of Developer Managed VMs | Number of developer managed VMs.<br>**NOTE**<br>This is published if the cluster is Workload Management enabled or there are developer managed VMs under the cluster.<br>Key: summary\|numberDeveloperManagedVMs |
| Namespaces\|Config Status | Workload Management configuration status.<br>**NOTE**<br>This is published if the cluster is Workload Management enabled.<br>Key: namespaces\|configStatus |
| Namespaces\|Kubernetes Status | Kubernetes status.<br>**NOTE**<br>This is published if the cluster is Workload Management enabled.<br>Key: namespaces\|kuberntesStatus |

**Reclaimable Metrics for Cluster Compute Resources**

Reclaimable metrics provide information about reclaimable resources.

| Metric Name | Description |
|---|---|
| Idle VMs\|CPU (vCPUs) | Number of reclaimable vCPUs of Idle VMs within the cluster.<br>Key: reclaimable\|idle_vms\|cpu |
| Idle VMs\|Disk Space (GB) | Reclaimable disk space of Idle VMs within the cluster.<br>Key: reclaimable\|idle_vms\|disksapce |
| Idle VMs\|Memory (KB) | Reclaimable memory of Idle VMs within the cluster.<br>Key: reclaimable\|idle_vms\|mem |
| Idle VMs\|Potential Savings | Potential saving after reclamation of resources of Idle VMs within the cluster.<br>Key: reclaimable\|idle_vms\|cost |
| Powered Off VMs\|Disk Space (GB) | Reclaimable disk space of Powered Off VMs within the cluster.<br>Key: reclaimable\|poweredOff_vms\| diskspace |
| Powered Off VMs\|Potential Savings | Potential saving after reclamation of resources of Powered Off VMs within the cluster.<br>Key: reclaimable\|poweredOff_vms\|cost |
| VM Snapshots\|Disk Space (GB) | Reclaimable disk space of VM Snapshots within the cluster.<br>Key: reclaimable\| vm_snapshots \| diskspace |
| VM Snapshots \|Potential Savings | Potential saving after reclamation of VM Snapshots within the cluster.<br>Key: reclaimable\| vm_snapshots \|cost |

**Cost Metrics for Cluster Compute Resources**

Cost metrics provide information about the cost.

| Metric Name | Description |
|---|---|
| Cluster CPU Base Rate | Base rate for Cluster CPU calculated by dividing the monthly total cluster CPU cost by cluster CPU utilization % and CPU cluster capacity (gHZ).<br>Key:cost\|cpuBaseRate |
| Cluster CPU Utilization (%) | Expected CPU utilization that is set by the user in cluster cost page.<br>Key:cost\|cpuExpectedUtilizationPct |
| Cluster Memory Base Rate | Cluster memory base rate calculated by dividing the monthly total cluster memory cost by cluster memory utilization % and memory cluster capacity (GB).<br>Key: cost\|memoryBaseRate |
| Cluster Memory Utilization (%) | Expected memory utilization that is set by the user in cluster cost page.<br>Key: cost\|memoryExpectedUtilizationPct |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Monthly Cluster Allocated Cost | Monthly cluster allocated cost calculated by subtracting the monthly cluster unallocated cost from the monthly cluster total cost.<br>Key: cost\|allocatedCost |
| Monthly Cluster Total Cost | Fully loaded compute cost of all hosts underneath the cluster.<br>Key: cost\|totalCost |
| Monthly Cluster Unallocated Cost | Monthly cluster unallocated cost calculated by subtracting the monthly cluster allocated cost from the monthly cluster total cost.<br>Key: cost\|unAllocatedCost |
| Monthly Total Cluster CPU Cost | Cost attributed to the cluster CPU from monthly cluster total cost.<br>Key: cost\|totalCpuCost |
| Monthly Total Cluster Memory Cost | Cost attributed to the cluster memory from monthly cluster total cost.<br>Key: cost\|totalMemoryCost |
| MTD Cluster CPU Utilization (GHz) | Month to date CPU utilization of the cluster.<br>Key: cost\|cpuActualUtilizationGHz |
| MTD Cluster Memory Utilization (GB) | Month to date memory utilization of the cluster.<br>Key: cost\|memoryActualUtilizationGB |
| Monthly Cluster Allocated Cost (Currency) | The monthly allocated cost of all VMs in a cluster.<br>cost\|clusterAllocatedCost |
| Cost\|Allocation\|Monthly Cluster Unallocated Cost (Currency) | The monthly unallocated is calculated by subtracting the monthly allocated cost from the cluster's cost.<br>cost\|clusterUnAllocatedCost |
| Aggregated Daily Total Cost | Daily aggregate daily total cost of the deleted VM present in the host system.<br>Key: Cost\|aggregatedDailyTotalCost |
| Aggregated Deleted VM Daily Total Cost | Daily aggregate cost of the deleted VM present in the host system.<br>Key: Cost\|aggregatedDeletedVmDailyTotalCost |

## Profiles Metrics for Cluster Compute Resources

Profiles metrics provide information about the profile specific capacity.

| Metric Name | Description |
|---|---|
| Profiles\|Capacity Remaining Profile (Average) | The capacity remaining in terms of fitting the average consumer.<br>Key: Profiles\|capacityRemainingProfile_<profile uuid> |
| Profiles\|Capacity Remaining Profile (<custom profile name>) | Published for custom profiles enabled from policy on Cluster Compute Resource.<br>Key: Profiles\|capacityRemainingProfile_<profile uuid> |

### Capacity Allocation Metrics for Cluster Compute Resources

Capacity allocation metrics provide information about the allotment of capacity, see Capacity Analytics Generated Metrics.

### Virtual Machine Operations Metrics for Clusters

VM operations metrics provide information about the actions performed on VMs. The following are some important points you must know about VM operation metrics for clusters.

- VM operations metrics is not collected for custom data centers.
- If you edit a VM settings and do not perform any action, still it is considered as VM reconfigure operation.
- During Revert Snapshot, VMs are powered-off, but this operation is not counted under VM Power-off metric.
- Adding ESXi with VMs is not counted under VM Create metric.
- Removing ESXi with VMs is not coutned under VM Remove metric.
- VM hardstop operation is not counted under VM Power Off metric.

| Metric Name | Description |
|---|---|
| Inventory | |
| VM Clone | This metric displays the number of clone operations on the virtual machine.<br>Key: Inventory\|VM Clone |
| VM Create | This metric displays the number of create operations on the virtual machine.<br>Key: Inventory\|VM Create |
| VM Delete | This metric displays the number of delete operations on the virtual machine.<br>Key: Inventory\|VM Delete |
| VM Reconfigure | This metric displays the number of reconfigure operations on the virtual machine.<br>Key: Inventory\|VM Reconfigure |
| VM Register | This metric displays the number of register operations on the virtual machine.<br>Key: Inventory\|VM Register |
| VM Template Deploy | This metric displays the number templates deployed on the virtual machine.<br>Key: Inventory\|VM Template Deploy |
| VM Unregister | This metric displays the number of unregister operations on the virtual machine.<br>Key: Inventory\|VM Unregister |
| Location | |
| Storage vMotion | This metric displays the number of migrations with vMotion (datastore change operations for Powered-on VMs).<br>Key: Location\|Storage vMotion |
| VM Datastore Change (powered-off VMs) | This metric displays the number of datastore change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-off VMs) | This metric displays the number of host and datastore change operations, for powered-off and suspended virtual machines. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: Location\|VM Host and Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-on VMs) | This metric displays the number of host and datastore change operations, for powered-on and suspended virtual machines.<br>Key: Location\|VM Host and Datastore Change (powered-on VMs) |
| VM Host Change (powered-off VMs) | This metric displays the number of host change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Host Change (powered-off VMs) |
| vMotion | This metric displays the number of migrations with vMotion (host change operations for powered-on VMs).<br>Key: Location\|vMotion |
| State | |
| VM Guest Reboot | This metric displays the number of reboot operations on the virtual machine guest.<br>Key: State\|VM Guest Reboot |
| VM Guest Shutdown | This metric displays the number of shutdown operations on the virtual machine guest.<br>Key: State\|VM Guest Shutdown |
| VM Power Off | This metric displays the number of power-off operations on the virtual machine.<br>Key: State\|VM Power Off |
| VM Power On | This metric displays the number of power-on operations on the virtual machine.<br>Key: State\|VM Power On |
| VM Reset | This metric displays the number of reset operations on the virtual machine guest.<br>Key: State\|VM Reset |
| VM Standby Guest | This metric displays the number of standby operations on the virtual machine guest.<br>Key: State\|VM Standby Guest |
| VM Suspend | This metric displays the number of suspend operations on the virtual machine.<br>Key: State\|VM Suspend |

**Disabled Metrics**

The following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

You can enable these metrics in the Policy workspace. For more information, see *Metrics and Properties Details*.

| Metric Name | Key |
|---|---|
| CPU\|Capacity Available to VMs (mhz) | cpu\|totalCapacity_average |
| CPU\|IO Wait (msec) | cpu\|iowait |
| CPU\|Reserved Capacity (mhz) | cpu\|reservedCapacity_average |
| CPU\|Total Wait (msec) | cpu\|wait |
| Datastore I/O\|Max Observed Number of Outstanding IO Operations | datastore\|maxObserved_OIO |
| Datastore I/O\|Max Observed Read Rate (kbps) | datastore\|maxObserved_Read |
| Datastore I/O\|Max Observed Reads per second | datastore\|maxObserved_NumberRead |
| Datastore I/O\|Max Observed Write Rate (kbps) | datastore\|maxObserved_Write |
| Datastore I/O\|Max Observed Writes per second | datastore\|maxObserved_NumberWrite |
| Storage\|Total Usage (kbps) | storage\|usage_average |
| Summary\|Average Provisioned Capacity per Running VM (mhz) | summary\|avg_vm_cpu |
| Summary\|Average Provisioned Memory per Running VM (kb) | summary\|avg_vm_mem |
| Summary\|Average Provisioned Memory per Running VM (kb) | summary\|avg_vm_mem |
| Summary\|Maximum Number of VMs | summary\|max_number_vms |
| Summary\|Workload Indicator | summary\|workload_indicator |
| Network I/O\|Max Observed Received Throughput (KBps) | net\|maxObserved_Rx_KBps |
| Network I/O\|Max Observed Throughput (KBps) | net\|maxObserved_KBps |
| Network I/O\|Max Observed Transmitted Throughput (KBps) | net\|maxObserved_Tx_KBps |
| Diskspace\|Not Shared (GB) | Space used by VMs that is not shared.<br>Key: diskspace\|notshared |

### Resource Pool Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, CPU usage, memory, and summary metrics for resource pool objects.

Resource Pool metrics include capacity and badge metrics. See definitions in:
- Capacity Analytics Generated Metrics
- Badge Metrics

#### Configuration Metrics for Resource Pools

Configuration metrics provide information about memory and CPU allocation configuration.

| Metric Name | Description |
|---|---|
| Memory Allocation Reservation | Memory Allocation Reservation.<br>Key: config\|mem_alloc_reservation |

#### CPU Usage Metrics for Resource Pools

CPU usage metrics provide information about CPU use.

| Metric Name | Description |
|---|---|
| Capacity Demand Entitlement (%) | CPU Capacity Demand Entitlement Percentage.<br>Key: cpu\|capacity_demandEntitlementPct |
| Capacity entitlement (MHz) | CPU Capacity Entitlement.<br>Key: cpu\|capacity_entitlement |
| CPU Contention (%) | CPU capacity contention.<br>Key: cpu\|capacity_contentionPct |
| Demand (MHz) | CPU demand in megahertz.<br>Key: cpu\|demandmhz |
| Overall CPU Contention | Overall CPU contention in milliseconds.<br>Key: cpu\|capacity_contention |
| Usage | Average CPU use in megahertz.<br>Key: cpu\|usagemhz_average |
| Effective limit | CPU effective limit.<br>Key: cpu\|effective_limit |
| Reservation Used | CPU reservation used.<br>Key: cpu\|reservation_used |
| Estimated entitlement | CPU estimated entitlement.<br>Key: cpu\|estimated_entitlement |
| Dynamic entitlement | CPU dynamic entitlement.<br>Key: cpu\|dynamic_entitlement |
| Demand without overhead | Value of demand excluding any overhead.<br>Key: cpu\|demand_without_overhead |

**Memory Metrics for Resource Pools**

Memory metrics provide information about memory use and allocation.

| Metric Name | Description |
|---|---|
| Balloon | Amount of memory currently used by the virtual machine memory control.<br>Key: mem\|vmmemctl_average |
| Compression Rate | Compression rate in kilobytes per second.<br>Key: mem\|compressionRate_average |
| Consumed | Amount of host memory consumed by the virtual machine for guest memory.<br>Key: mem\|consumed_average |
| Contention | Machine contention.<br>Key: mem\|host_contentionPct |
| Guest usage | Guest memory entitlement.<br>Key: mem\|guest_usage |
| Guest demand | Guest memory entitlement.<br>Key: mem\|guest_demand |
| Contention (KB) | Machine contention in kilobytes.<br>Key: mem\|host_contention |
| Decompression Rate | Decompression rate in kilobytes per second. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: mem\|decompressionRate_average |
| Granted | Average of memory available for use.<br>Key: mem\|granted_average |
| Guest Active | Amount of memory that is actively used.<br>Key: mem\|active_average |
| VM Overhead | Memory overhead reported by host.<br>Key: mem\|overhead_average |
| Shared | Amount of shared memory.<br>Key: mem\|shared_average |
| Reservation Used | Memory Reservation Used.<br>Key: mem\|reservation_used |
| Dynamic Entitlement | Memory Dynamic Entitlement.<br>Key: mem\|dynamic_entitlement |
| Effective Limit | Memory Effective Limit.<br>Key: mem\|effective_limit |
| Swap In Rate | Rate at which memory is swapped from disk into active memory during the interval.<br>Key: mem\|swapinRate_average |
| Swap Out Rate | Rate at which memory is being swapped from active memory to disk during the current interval.<br>Key: mem\|swapoutRate_average |
| Swapped | Amount of unreserved memory.<br>Key: mem\|swapped_average |
| Usage (%) | Memory currently in use as a percentage of total available memory.<br>Key: mem\|usage_average |
| Zero | Amount of memory that is all zero.<br>Key: mem\|zero_average |
| Zipped (KB) | Latest zipped memory in kilobytes.<br>Key: mem\|zipped_latest |
| Swap In (KB) | Amount of memory swapped in kilobytes.<br>Key: mem\|swapin_average |
| Swap Out (KB) | Amount of memory swapped out in kilobytes.<br>Key: mem\|swapout_average |
| Swap Used | Amount of memory used for swap space in kilobytes.<br>Key: mem\|swapused_average |
| Total Capacity | Total capacity.<br>Key: mem\|guest_provisioned |

## Summary Metrics for Resource Pools

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| Number of Running VMs | Number of running virtual machines.<br>Key: summary\|number_running_vms |
| Total Number of VMs | Total number of virtual machines.<br>**NOTE**<br>This shows the total number of VMs excluding VM templates.<br>Key: summary\|total_number_vms |
| IO Wait (ms) | IO wait time in milliseconds.<br>Key: summary\|iowait |
| Number of VM Templates | Number of VM Templates.<br>Key: summary\|number_vm_templates |

### Data Center Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects CPU usage, disk, memory, network, storage, disk space, and summary metrics for data center objects.

Data center metrics include capacity and badge metrics. See definitions in:
- Capacity Analytics Generated Metrics
- Badge Metrics

### Data Center Metrics for ROI Dashboard

Data center metrics provide information about data center savings across vCenters.

| Metric Name | Description |
|---|---|
| Realized Cost Savings | |
| Realized Savings Idle Cost | This metric displays the total realize savings for VMs across all vCenters.<br>Key: cost\|realized_savings\|realizedIdleCost |
| Realized Savings Powered Off Cost (AOA) | This metric displays the total realized savings for powered off VMs across all vCenters.<br>Key: cost\|realized_savings\|realizedPoweredOffCost |
| Realized Savings Snapshot Space Cost (AOA) | This metric displays the snapshots space saved across all vCenters.<br>Key: cost\|realized_savings\|realizedSnapshotSpaceCost |
| Realized Savings Oversized Cost (AOA) | This metric displays the oversized savings across all vCenters.<br>Key: cost\|realized_savings\|realizedOversizedCost |
| Realized Savings Orphaned Disk Space Cost (AOA) | This metric displays the amount of disk space saved by orphaned disks across all vCenters.<br>Key: cost\|realized_savings\|realizedOrphanedDiskSpaceCost |
| Realized Savings Reclaimable Host Cost (AOA) | This metric displays the amount of reclaimable host savings across all vCenters.<br>Key: cost\|realized_savings\|realizedReclaimableHostCost |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Realized vCPUs from Oversized VMs | This metric displays the number of vCPUs realized across all vCenters.<br>Key: realized\|realizedVCpus |
| Compute Realized Memory from Oversized VMs | This metric displays the amount of memory realized from oversized VMs across all vCenters.<br>Key: compute_realized\|realizedOversizedMem |
| Realized Potential Memory Consumed from Oversized VMs | This metric displays the potential memory consumed from oversized VMs across all vCenters.<br>Key: realized\|realizedPotentialMemConsumed |
| Compute Realized vCPUs from Oversized VMs | This metric displays the realized vCPUs from oversized VMs across all vCenters.<br>Key: compute_realized\|realizedOversizedVCpus |
| Compute Realized vCPUs from Idle VMs | This metric displays the realized vCPUs from idle VMs across all vCenters.<br>Key: compute_realized\|realizedIdleVCpus |
| Compute Realized Memory from Idle VMs | This metric displays the amount of memory realized from idle VMs across all vCenters.<br>Key: compute_realized\|realizedIdleMem |
| Disk Space Realized Idle VMs | This metric displays the amount of disk space realized from idle VMs across all vCenters.<br>Key: storage_realized\|realizedIdleDiskSpace |
| Disk Space Realized PoweredOff VMs | This metric displays the amount of disk space realized from powered off VMs across all vCenters.<br>Key: storage_realized\|realizedPoweredOffDiskSpace |
| Disk Space Realized VM Snapshots | This metric displays the amount of disk space realized from VM snapshots across all vCenters.<br>Key: storage_realized\|realizedSnapshotSpace |
| Disk Space Realized Orphaned Disks | This metric displays the amount of disk space realized from orphaned disks across all vCenters.<br>Key: storage_realized\|realizedIdleDiskSpace |
| Realized Savings Total Realized Cost | This metric displays the total realized cost across all vCenters.<br>Key: cost\|realized_savings\|realizedTotalCost |

**CPU Usage Metrics for Data Centers**

CPU usage metrics provide information about CPU use.

| Metric Name | Description |
|---|---|
| Capacity Usage (%) | Percent capacity used.<br>Key: cpu\|capacity_usagepct_average |
| CPU Contention (%) | CPU capacity contention.<br>Key: cpu\|capacity_contentionPct |
| Demand (%) | CPU demand percentage.<br>Key: cpu\|demandPct |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Demand | Demand in megahertz.<br>Key: cpu\|demandmhz |
| Demand (MHz) | CPU utilization level based on descendant virtual machines utilization. This Includes reservations, limits, and overhead to run the virtual machines.<br>Key: cpu\|demandmhz |
| Overhead (KB) | Amount of CPU overhead.<br>Key: cpu\|overhead_average |
| Demand without overhead | Value of demand excluding any overhead.<br>Key: cpu\|demand_without_overhead |
| Total Wait | CPU time spent on idle state.<br>Key: cpu\|wait |
| Number of CPU Sockets | Number of CPU sockets.<br>Key: cpu\|numpackages |
| Overall CPU Contention (ms) | Overall CPU contention in milliseconds.<br>Key: cpu\|capacity_contention |
| Host Provisioned Capacity (MHz) | Host provisioned capacity in megahertz.<br>Key: cpu\|capacity_provisioned |
| Provisioned vCPU(s) | Provisioned vCPU(s).<br>Key: cpu\|corecount_provisioned |
| Reserved Capacity (MHz) | The sum of the reservation properties of the (immediate) children of the host's root resource pool.<br>Key: cpu\|reservedCapacity_average |
| Usage | Average CPU usage in megahertz.<br>Key: cpu\|usagemhz_average |
| IO Wait | IO wait time in milliseconds.<br>Key: cpu\|iowait |
| Provisioned Capacity | Provisioned Capacity.<br>Key: cpu\|vm_capacity_provisioned |
| Stress Balance Factor | Stress Balance Factor.<br>Key: cpu\|stress_balance_factor |
| Lowest Provider Capacity Remaining | Lowest Provider Capacity Remaining.<br>Key: cpu\|min_host_capacity_remaining |
| Workload Balance Factor | Workload Balance Factor.<br>Key: cpu\|workload_balance_factor |
| Highest Provider Workload | Highest Provider Workload.<br>Key: cpu\|max_host_workload |
| Host workload Max-Min Disparity | Difference of Max and Min host workload in the container.<br>Key: cpu\|host_workload_disparity |
| Host stress Max-Min Disparity | Difference of Max and Min host stress in the container.<br>Key: cpu\|host_stress_disparity |
| Total Capacity (MHz) | Total CPU resources configured on the descendant ESXi hosts.<br>Key: cpu\|capacity_provisioned |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Usable Capacity (MHz) | The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services.<br>Key: cpu\|haTotalCapacity_average |

**Disk Metrics for Data Centers**

Disk metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: disk\|commandsAveraged_average |
| Total Latency (ms) | Average amount of time taken for a command from the perspective of the guest operating system. This metric is the sum of the Kernel Latency and Physical Device Latency metrics.<br>Key: disk\|totalLatency_average |
| Total Throughput (KBps) | Average of the sum of the data read and written for all the disk instances of the host or virtual machine.<br>Key: disk\|usage_average |
| Total queued outstanding operations | Sum of queued operations and outstanding operations.<br>Key: disk\|sum_queued_oio |
| Max observed OIO | Max observed IO for a disk.<br>Key: disk\|max_observed |

**Memory Metrics for Data Centers**

Memory metrics provide information about memory use and allocation.

| Metric Name | Description |
|---|---|
| Contention (%) | Machine Contention Percentage.<br>Key: mem\|host_contentionPct |
| Machine Demand (KB) | Memory machine demand in kilobytes.<br>Key: mem\|host_demand |
| ESX System Usage | Memory usage by the VM kernel and ESX user-level services.<br>Key: mem\|host_systemUsage |
| Provisioned Memory (KB) | Provisioned host memory in kilobytes.<br>Key: mem\|host_provisioned |
| Reserved Capacity (KB) | Reserved memory capacity in kilobytes.<br>Key: mem\|reservedCapacity_average |
| Usable Memory (KB) | Usable host memory in kilobytes.<br>Key: mem\|host_usable |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Host Usage | Host memory use in kilobytes.<br>Key: mem\|host_usage |
| Usage/Usable (%) | Percent host memory used.<br>Key: mem\|host_usagePct |
| VM Overhead | Memory overhead reported by host.<br>Key: mem\|overhead_average |
| Stress Balance Factor | Stress Balance Factor.<br>Key: mem\|stress_balance_factor |
| Lowest Provider Capacity Remaining | Lowest Provider Capacity Remaining.<br>Key: mem\|min_host_capacity_remaining |
| Workload Balance Factor | Workload Balance Factor.<br>Key: mem\|workload_balance_factor |
| Highest Provider Workload | Highest Provider Workload.<br>Key: mem\|max_host_workload |
| Host workload Max-Min Disparity | Difference of Max and Min host workload in the container.<br>Key: mem\|host_workload_disparity |
| Host stress Max-Min Disparity | Difference of Max and Min host stress in the container.<br>Key: mem\|host_stress_disparity |
| Utilization (KB) | Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines.<br>Key: mem\|total_need |
| Total Capacity (KB) | Total physical memory configured on descendant ESXi hosts.<br>Key: mem\|host_provisioned |
| Usable Capacity (KB) | The usable memory resources available for the virtual machines after considering reservations for vSphere HA and other vSphere services.<br>Key: mem\|haTotalCapacity_average |

## Network Metrics for Data Centers

Network metrics provide information about network performance.

| Metric Name | Description |
|---|---|
| Packets Dropped | Percentage of packets dropped.<br>Key: net\|droppedPct |
| Max Observed Throughput | Max observed rate of network throughput.<br>Key: net\|maxObservedKBps |
| Data Transmit Rate | Average amount of data transmitted per second.<br>Key: net\|transmitted_average |
| Data Receive Rate | Average amount of data received per second.<br>Key: net\|received_average |
| Total Throughput (KBps) | The sum of the data transmitted and received for all the NIC instances of the host or virtual machine. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: net\|usage_average |

**Storage Metrics for Data Centers**

Storage metrics provide information about storage use.

| Metric Name | Description |
|---|---|
| Total Usage | Total throughput rate.<br>Key: storage\|usage_average |

**Datastore Metrics for Data Centers**

Datastore metrics provide information about Datastore use.

| Metric Name | Description |
|---|---|
| Outstanding IO requests | OIO for datastore.<br>Key: datastore\|demand_oio |
| Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: datastore\|numberReadAveraged_average |
| Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: datastore\|numberWriteAveraged_average |
| Read Throughput (KBps) | Amount of data read in the performance interval.<br>Key: datastore\|read_average |
| Write Throughput (KBps) | Amount of data written to disk in the performance interval.<br>Key: datastore\|write_average |

**Disk Space Metrics for Data Centers**

Disk space metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| Virtual machine used | Used virtual machine disk space in gigabytes.<br>Key: diskspace\|used |
| Total disk space used | Total disk space used on all datastores visible to this object.<br>Key: diskspace\|total_usage |
| Total disk space | Total disk space on all datastores visible to this object.<br>Key: diskspace\|total_capacity |
| Total provisioned disk space | Total provisioned disk space on all datastores visible to this object.<br>Key: diskspace\|total_provisioned |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Shared Used (GB) | Shared disk space in gigabytes.<br>Key: diskspace\|shared |
| Snapshot Space (GB) | Snapshot disk space in gigabytes.<br>Key: diskspace\|snapshot |
| Virtual Disk Used (GB) | Used virtual disk space in gigabytes.<br>Key: diskspace\|diskused |
| Number of Virtual Disks | Number of Virtual Disks.<br>Key: diskspace\|numvmdisk |
| Utilization (GB) | Storage space used on the connected vSphere Datastores.<br>Key: diskspace\|total_usage |
| Total Capacity (GB) | Total storage space available on the connected vSphere datastores.<br>Key: diskspace\|total_capacity |

**Summary Metrics for Data Centers**

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| Number of Running Hosts | Number of hosts that are ON.<br>Key: summary\|number_running_hosts |
| Number of Running VMs | Number of running virtual machines.<br>Key: summary\|number_running_vms |
| Maximum Number of VMs | Maximum number of virtual machines.<br>Key: summary\|max_number_vms |
| Number of Clusters | Total number of clusters.<br>Key: summary\|total_number_clusters |
| Number of Hosts | Total number of hosts.<br>Key: summary\|total_number_hosts |
| Number of VMs | Total number of virtual machines.<br>Key: summary\|total_number_vms |
| Total Number of Datastores | Total number of datastores.<br>Key: summary\|total_number_datastores |
| Number of VCPUs on Powered On VMs | Total number of VCPUs of virtual machines that are powered on.<br>Key: summary\|number_running_vcpus |
| Workload Indicator | Workload indicator.<br>Key: summary\|workload_indicator |
| Average Running VM Count per Running Host | Average number of running virtual machines per running host.<br>Key: summary\|avg_vm_density |
| WLP | Displays the VM migration trend as part of workload optimization. These metrics are deactivated by default. You must activate them from policies. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | • Fail Count: Number of failed VM move attempts in the last daily cycle.<br>• Number of runs: The total number of times the WLP was run during the last daily cycle.<br>• Success Count: The number of successful VM moves during the last daily cycle. |

**Reclaimable Metrics for Data Centers**

Reclaimable metrics provide information about reclaimable resources.

| Metric Name | Description |
|---|---|
| CPU (vCPUs) | Number of reclaimable vCPUs within the data center.<br>Key: reclaimable\|cpu |
| Disk Space | Reclaimable disk space within the data center.<br>Key: reclaimable\|diskspace |
| Potential Savings | Potential saving after reclamation of resources of all reclaimable VMs (Idle VMs, Powered Off VMs, VM snapshots) within the data center.<br>Key: reclaimable\|cost |
| Memory (KB) | Reclaimable memory within the data center.<br>Key: reclaimable\|mem |
| Virtual Machines | Number of VMs having reclaimable resources (Memory, disk space, vCPU) within the data center.<br>Key: reclaimable\|vm_count |
| Idle VMs\|Potential Savings | Potential saving after reclamation of resources of Idle VMs within the data center.<br>Key: reclaimable\|idle_vms\|cost |
| Powered Off VMs\|Potential Savings | Potential saving after reclamation of resources of Powered Off VMs within the data center.<br>Key: reclaimable\|poweredOff_vms\|cost |
| VM Snapshots\|Potential Savings | Potential saving after reclamation of VM snapshots within the data center.<br>Key: reclaimable\|vm_snapshots \|cost |
| Reclaimable\|Orphaned Disks\|Potential Savings (Currency) | Displays the potential savings after reclaimation of disk space by removing orphaned VMDks from all datastores under datacenter.<br>reclaimable\|cost |
| Reclaimable\|Number of Orphaned Disks | Number of reclaimable orphaned disks is the sum of all orphaned disks on it's datastore.<br>reclaimable\|orphaned_disk_count |
| Reclaimable Host Cost | This metric calculates the cumulated reclaimable host cost at cluster level and displays the same metric at data center level.<br>Reclaimable\|Reclaimable Host Cost (Currency) |

## Cost Metrics for Data Centers

Cost metrics provide information about the cost.

**NOTE**
The Effective metrics based on the capacity model assigned at the cluster (Demand/Allocation) pick the corresponding metric from underneath the clusters.

| Metric Name | Description |
|---|---|
| Monthly Cluster Aggregated Allocated Cost | This metric displays the sum of the monthly allocated cost for both cluster and unclustered hosts.<br>Key: Cost\|Allocation\|Monthly Cluster Aggregated Allocated Cost |
| Monthly Cluster Aggregated Unallocated Cost | This metric displays the sum of both cluster and unclustered hosts unallocated cost.<br>Key: Cost\|Allocation\|Monthly Cluster Aggregated Unallocated Cost |
| Monthly Datastore Aggregated Allocated Cost | This metric displays the monthly aggregated allocated cost at datastore level.<br>Key: Cost\|Allocation\|Monthly Datastore Aggregated Allocated Cost |
| Monthly Datastore Aggregated Unallocated Cost | This metric displays the monthly aggregated unallocated cost at datastore level.<br>Key: Cost\|Allocation\|Monthly Datastore Aggregated Unallocated Cost |
| Monthly Cluster Effective Aggregated Allocated Cost | This metric displays the Monthly Cluster Effective Aggregated Allocated Cost cost.<br>Key: Cost\|Monthly Cluster Effective Aggregated Allocated Cost |
| Monthly Cluster Effective Aggregated Unallocated Cost | This metric displays the Monthly Cluster Effective Aggregated Unallocated Cost.<br>Key: Cost\|Monthly Cluster Effective Aggregated Unallocated Cost |
| Monthly Datastore Effective Aggregated Allocated Cost | This metric displays the Monthly Datastore Effective Aggregated Allocated Cost at datastore level.<br>Key: Cost\|Allocation\|Monthly Datastore Effective Aggregated Allocated Cost |
| Monthly Datastore Effective Aggregated Unallocated Cost | This metric displays the Monthly Datastore Effective Aggregated Unallocated Cost at datastore level.<br>Key: Cost\|Allocation\|Monthly Datastore Effective Aggregated Unallocated Cost |
| **NOTE**<br>The eight allocation metrics mentioned above, helps you to calculate the cost when you enable the allocation model for costing. The same set of metrics are available when you enable demand model also. | |
| Monthly Cluster Aggregated Cost | This metric displays the sum of monthly aggregated allocated and unallocated cost for clusters.<br>Key: cost\|clusterCost |
| Monthly Cluster Aggregated Unallocated Cost | This metric displays the sum of the monthly unallocated cost for both cluster and unclustered hosts.<br>Key: Cost\|Monthly Cluster Aggregated Unallocated Cost |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Monthly Datacenter Aggregated Total Cost | Monthly aggregated total cost for the data center.<br>Key: Cost\|Monthly Datacenter Aggregated Total Cost |
| Monthly Datastore Total Cost | Monthly data store total cost.<br>Key: cost\|totalCost |
| Monthly Datastore Aggregated Allocated Cost | Monthly aggregated allocated cost for the datastore.<br>Key: cost\|aggrDataStoreAllocatedCost |
| Monthly Datastore Aggregated Unallocated Cost | Monthly aggregated unallocated cost for the datastore.<br>Key: cost\|aggrDataStoreUnallocatedCost |
| Monthly VM Aggregated Direct Cost | Month to date aggregated VM direct cost across all the VMs under the data center.<br>Key: cost\|vmDirectCost |

**Virtual Machine Operations Metrics for Data Centers**

VM operations metrics provide information about the actions performed on the VMs in the datacenter. The following are some important points you must know about VM operation metrics for data centers.

- VM operations metrics is not collected for custom data centers.
- If you edit a VM settings and do not perform any action, still it is considered as VM reconfigure operation.
- During Revert Snapshot, VMs are powered-off, but this operation is not counted under VM Power-off metric.
- Adding ESXi with VMs is not counted under VM Create metric.
- Removing ESXi with VMs is not coutned under VM Remove metric.
- VM hardstop operation is not counted under VM Power Off metric.

| Metric Name | Description |
|---|---|
| Inventory | |
| VM Clone | This metric displays the number of clone operations on the virtual machine.<br>Key: Inventory\|VM Clone |
| VM Create | This metric displays the number of create operations on the virtual machine.<br>Key: Inventory\|VM Create |
| VM Delete | This metric displays the number of delete operations on the virtual machine.<br>Key: Inventory\|VM Delete |
| VM Reconfigure | This metric displays the number of reconfigure operations on the virtual machine.<br>Key: Inventory\|VM Reconfigure |
| VM Register | This metric displays the number of register operations on the virtual machine.<br>Key: Inventory\|VM Register |
| VM Template Deploy | This metric displays the number templates deployed on the virtual machine.<br>Key: Inventory\|VM Template Deploy |
| VM Unregister | This metric displays the number of unregister operations on the virtual machine. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: Inventory\|VM Unregister |
| Location | |
| Storage vMotion | This metric displays the number of migrations with vMotion (datastore change operations for Powered-on VMs).<br>Key: Location\|Storage vMotion |
| VM Datastore Change (powered-off VMs) | This metric displays the number of datastore change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-off VMs) | This metric displays the number of host and datastore change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Host and Datastore Change (powered-off VMs) |
| VM Host and Datastore Change (powered-on VMs) | This metric displays the number of host and datastore change operations, for powered-on and suspended virtual machines.<br>Key: Location\|VM Host and Datastore Change (powered-on VMs) |
| VM Host Change (powered-off VMs) | This metric displays the number of host change operations, for powered-off and suspended virtual machines.<br>Key: Location\|VM Host Change (powered-off VMs) |
| vMotion | This metric displays the number of migrations with vMotion (host change operations for powered-on VMs).<br>Key: Location\|vMotion |
| State | |
| VM Guest Reboot | This metric displays the number of reboot operations on the virtual machine guest.<br>Key: State\|VM Guest Reboot |
| VM Guest Shutdown | This metric displays the number of shutdown operations on the virtual machine guest.<br>Key: State\|VM Guest Shutdown |
| VM Power Off | This metric displays the number of power-off operations on the virtual machine.<br>Key: State\|VM Power Off |
| VM Power On | This metric displays the number of power-on operations on the virtual machine.<br>Key: State\|VM Power On |
| VM Reset | This metric displays the number of reset operations on the virtual machine guest.<br>Key: State\|VM Reset |
| VM Standby Guest | This metric displays the number of standby operations on the virtual machine guest.<br>Key: State\|VM Standby Guest |
| VM Suspend | This metric displays the number of suspend operations on the virtual machine.<br>Key: State\|VM Suspend |

## Disabled Metrics

The following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

You can enable these metrics in the Policy workspace. For more information, see Metrics and Properties Details.

| Metric Name | Key |
|---|---|
| Datastore I/O\|Max Observed Number of Outstanding IO Operations (IOPS) | datastore\|maxObserved_OIO |
| Datastore I/O\|Max Observed Read Rate (KBps) | datastore\|maxObserved_Read |
| Datastore I/O\|Max Observed Reads per second (IOPS) | datastore\|maxObserved_NumberRead |
| Datastore I/O\|Max Observed Write Rate (KBps) | datastore\|maxObserved_Write |
| Datastore I/O\|Max Observed Writes per second (IOPS)0 | datastore\|maxObserved_NumberWrite |
| Max Observed Transmitted Throughput | Max observed transmitted rate of network throughput.<br>Key: net\|maxObserved_Tx_KBps |
| Max Observed Received Throughput | Max observed received rate of network throughput.<br>Key: net\|maxObserved_Rx_KBps |
| Not Shared (GB) | Unshared disk space in gigabytes.<br>Key: diskspace\|notshared |

### Custom Data Center Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects CPU usage, memory, summary, network, and datastore metrics for custom data center objects.

Custom data center metrics include capacity and badge metrics. See definitions in:

- Capacity Analytics Generated Metrics
- Badge Metrics

## CPU Usage Metrics for Custom Data Centers

CPU usage metrics provide information about CPU use.

| Metric Name | Description |
|---|---|
| Host Provisioned Capacity | Host provisioned capacity (MHz).<br>Key: cpu\|capacity_provisioned |
| Provisioned vCPU(s) | Provisioned vCPU(s).<br>Key: cpu\|corecount_provisioned |
| Demand without overhead | Value of demand excluding any overhead.<br>Key: cpu\|demand_without_overhead |
| Number of hosts stressed | Number of hosts stressed.<br>Key: cpu\|num_hosts_stressed |
| Stress Balance Factor | Stress balance factor.<br>Key: cpu\|stress_balance_factor |
| Lowest Provider Capacity Remaining | Lowest provider capacity remaining. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: cpu\|min_host_capacity_remaining |
| Workload Balance Factor | Workload balance factor.<br>Key: cpu\|workload_balance_factor |
| Highest Provider Workload | Highest provider workload.<br>Key: cpu\|max_host_workload |
| Host workload Max-Min Disparity | Host workload max-min disparity.<br>Key: cpu\|host_workload_disparity |
| Host stress Max-Min Disparity | Difference of max and min host stress in the container.<br>Key: cpu\|host_stress_disparity |
| Demand (MHz) | CPU utilization level based on descendant virtual machines utilization. This Includes reservations, limits, and overhead to run the virtual machines.<br>Key: cpu\|demandmhz |
| Total Capacity (MHz) | Total CPU resources configured on the descendant ESXi hosts.<br>Key: cpu\|capacity_provisioned |
| Usable Capacity (MHz) | The usable CPU resources that are available for the virtual machines after considering reservations for vSphere High Availability (HA) and other vSphere services.<br>Key: cpu\|haTotalCapacity_average |

**Memory Metrics for Custom Data Centers**

Memory metrics provide information about memory use.

| Metric Name | Description |
|---|---|
| Usable Memory | Usable memory.<br>Key: mem\|host_usable |
| Machine Demand | Memory machine demand in KB.<br>Key: mem\|host_demand |
| Number of hosts stressed | Number of hosts stressed.<br>Key: mem\|num_hosts_stressed |
| Stress Balance Factor | Stress balance factor.<br>Key: mem\|stress_balance_factor |
| Lowest Provider Capacity Remaining | Lowest provider capacity remaining.<br>Key: mem\|min_host_capacity_remaining |
| Workload Balance Factor | Workload balance factor.<br>Key: mem\|workload_balance_factor |
| Highest Provider Workload | Highest provider workload.<br>Key: mem\|max_host_workload |
| Host workload Max-Min Disparity | Host workload max-min disparity.<br>Key: mem\|host_workload_disparity |
| Host stress max-min disparity | Host stress max-min disparity.<br>Key: mem\|host_stress_disparity |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Utilization (KB) | Memory utilization level based on the descendant virtual machines utilization. Includes reservations, limits, and overhead to run the Virtual Machines.<br>Key: mem\|total_need |
| Total Capacity (KB) | Total physical memory configured on descendant ESXi hosts.<br>Key: mem\|host_provisioned |
| Usable Capacity (KB) | The usable memory resources available for the virtual machines after considering reservations for vSphere HA and other vSphere services.<br>Key: mem\|haTotalCapacity_average |

**Summary Metrics for Custom Data Centers**

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| Number of Running VMs | Number of virtual machines that are ON.<br>Key: summary\|number_running_vms |
| Maximum Number of VMs | Maximum number of virtual machines.<br>Key: summary\|max_number_vms |
| Status | Status of the data center.<br>Key: summary\|status |
| WLP | Displays the VM migration trend as part of workload optimization. These metrics are deactivated by default. You must activate them from policies.<br>• Fail Count: Number of failed VM move attempts in the last daily cycle.<br>• Number of runs: The total number of times the WLP was run during the last daily cycle.<br>• Success Count: The number of successful VM moves during the last daily cycle. |

**Network Metrics for Custom Data Centers**

Network metrics provide information about network performance.

| Metric Name | Description |
|---|---|
| Usage Rate | The sum of the data transmitted and received for all the NIC instances of the host or virtual machine.<br>Key: net\|usage_average |
| Data Transmit Rate | Average amount of data transmitted per second.<br>Key: net\|transmitted_average |
| Data REceive Rate | Average amount of data received per second.<br>Key: net\|received_average |

### Datastore Metrics for Custom Data Centers

Datastore metrics provide information about datastore use.

| Metric Name | Description |
|---|---|
| Outstanding IO requests | OIO for datastore.<br>Key: datastore\|demand_oio |
| Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: datastore\|numberReadAveraged_average |
| Write IOPS | Average number of write commands issued per second during the collection interval.<br>Key: datastore\|numberWriteAveraged_average |
| Read Throughput (KBps) | Amount of data read in the performance interval.<br>Key: datastore\|read_average |
| Write Throughput (KBps) | Amount of data written to disk in the performance interval.<br>Key: datastore\|write_average |

### Reclaimable Metrics for Custom Data Centers

Reclaimable metrics provide information about reclaimable resources.

| Metric Name | Description |
|---|---|
| CPU (vCPUs) | Number of reclaimable vCPUs within the custom data center.<br>Key: reclaimable\|cpu |
| Disk Space | Reclaimable disk space within the custom data center.<br>Key: reclaimable\|diskspace |
| Potential Savings | Potential saving after reclamation of resources of all reclaimable VMs (Idle VMs, Powered Off VMs, VM snapshots) within the custom data center.<br>Key: reclaimable\|cost |
| Memory (KB) | Reclaimable memory within the custom data center.<br>Key: reclaimable\|mem |
| Number of Orphaned Disks | Number of reclaimable orphaned disks within the custom data center.<br>reclaimable\|orphaned_disk_count |
| Reclaimable\|Orphaned Disks\|Potential Savings | Potential savings in cost after reclamation of orphaned disks across the custom data center.<br>Key: reclaimable\|orphaned_disk\|cost<br><br>**NOTE**<br>The orphaned disk reclamation feature might not work as expected when VMware Aria OperationsV Mware Cloud Foundation Operations monitors multiple vCenters which use shared data stores. |
| Virtual Machines | Number of VMs having reclaimable resources (Memory, disk space, vCPU) within the custom data center. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: reclaimable\|vm_count |
| Idle VMs\|Potential Savings | Potential saving after reclamation of resources of Idle VMs within the custom data center.<br>Key: reclaimable\|idle_vms\|cost |
| Powered Off VMs\|Potential Savings | Potential saving after reclamation of resources of Powered Off VMs within the custom data center.<br>Key: reclaimable\|poweredOff_vms\|cost |
| VM Snapshots\|Potential Savings | Potential saving after reclamation of VM snapshots within the custom data center.<br>Key: reclaimable\|vm_snapshots \|cost |
| Reclaimable\|Orphaned Disks\|Potential Savings (Currency) | Displays the potential savings after reclaimation of disk space by removing orphaned VMDks from all datastores under custom datacenters.<br>reclaimable\|cost |
| Reclaimable\|Number of Orphaned Disks | Number of reclaimable orphaned disks is the sum of the numbers of orphaned disks on it's datastore.<br>reclaimable\|orphaned_disk_count |

**Disk Space Metrics for Custom Data Centers**

Disk space metrics provide information about disk use.

| Metric Name | Description |
|---|---|
| Utilization (GB) | Storage space used on the connected vSphere Datastores.<br>Key: diskspace\|total_usage |
| Total Capacity (GB) | Total storage space available on the connected vSphere datastores.<br>Key: diskspace\|total_capacity |

**Disabled Metrics**

The following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

You can enable these metrics in the Policy workspace. For more information, see Metrics and Properties Details.

| Metric Name | Key |
|---|---|
| Max Observed Throughput | Max observed rate of network throughput.<br>Key: net\|maxObserved_KBps |
| Max Observed Transmitted Throughput | Max observed transmitted rate of network throughput.<br>Key: net\|maxObserved_Tx_KBps |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Key |
|---|---|
| Max Observed Received Throughput | Max observed received rate of network throughput.<br>Key: net\|maxObserved_Rx_KBps |
| Max Observed Reads per second | Max observed average number of read commands issued per second during the collection interval.<br>Key: datastore\|maxObserved_NumberRead |
| Max Observed Read Rate | Max observed rate of reading data from the datastore.<br>Key: datastore\|maxObserved_Read |
| Max Observed Writes per second | Max observed average number of write commands issued per second during the collection interval.<br>Key: datastore\|maxObserved_NumberWrite |
| Max Observed Write Rate | Max observed rate of writing data from the datastore.<br>Key: datastore\|maxObserved_Write |
| Max Observed Number of Outstanding IO Operations | Max observer number of outstanding IO operations.<br>Key: datastore\|maxObserved_OIO |

### Storage Pod Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects datastore and disk space metrics for storage pod objects.

Storage Pod metrics include capacity and badge metrics. See definitions in:
- Capacity Analytics Generated Metrics
- Badge Metrics

**Table 408: Datastore Metrics for Storage Pods**

| Metric Name | Description |
|---|---|
| Read IOPS | Average number of read commands issued per second during the collection interval.<br>Key: datastore\|numberReadAveraged_average |
| Writes per second | Average number of write commands issued per second during the collection interval.<br>Key: datastore\|numberWriteAveraged_average |
| Read Throughput (KBps) | Amount of data read in the performance interval.<br>Key: datastore\|read_average |
| Write Throughput (KBps) | Amount of data written to disk in the performance interval.<br>Key: datastore\|write_average |
| Total Throughput (KBps) | Usage Average.<br>Key: datastore\|usage_average |
| Read Latency | Average amount of time for a read operation from the datastore.<br>Total latency = kernel latency + device latency.<br>Key: datastore\|totalReadLatency_average |
| Write Latency | Average amount of time for a write operation to the datastore.<br>Total latency = kernel latency + device latency.<br>Key: datastore\|totalWriteLatency_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Total Latency (ms) | The average amount of time taken for a command from the perspective of a Guest OS. This is the sum of Kernel Command Latency and Physical Device Command Latency.<br>Key: datastore\|totalLatency_average |
| Total IOPS | Average number of commands issued per second during the collection interval.<br>Key: datastore\|commandsAveraged_average |

**Table 409: Disk Space Metrics for Storage Pods**

| Metric Name | Description |
|---|---|
| Freespace | Unused space available on datastore.<br>Key: diskspace\|freespace |
| Total used | Total space used.<br>Key: diskspace\|disktotal |
| Capacity | Total capacity of datastore.<br>Key: diskspace\|capacity |
| Virtual Machine used | Space used by virtual machine files.<br>Key: diskspace\|used |
| Snapshot Space | Space used by snapshots.<br>Key: diskspace\|snapshot |

### VMware Distributed Virtual Switch Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects network and summary metrics for VMware distributed virtual switch objects.

VMware Distributed Virtual Switch metrics include badge metrics. See definitions in Badge Metrics.

**Table 410: Network Metrics for VMware Distributed Virtual Switches**

| Metric Name | Description |
|---|---|
| Total Ingress Traffic | Total ingress traffic (KBps).<br>Key: network\|port_statistics\|rx_bytes |
| Total Egress Traffic | Total egress traffic (KBps).<br>Key: network\|port_statistics\|tx_bytes |
| Egress Unicast Packets per second | Egress unicast packets per second.<br>Key: network\|port_statistics\|ucast_tx_pkts |
| Egress Multicast Packets per second | Egress multicast packets per second.<br>Key: network\|port_statistics\|mcast_tx_pkts |
| Egress Broadcast Packets per second | Egress broadcast packets per second.<br>Key: network\|port_statistics\|bcast_tx_pkts |
| Ingress Unicast Packets per second | Ingress unicast packets per second.<br>Key: network\|port_statistics\|ucast_rx_pkts |
| Ingress Multicast Packets per second | Ingress multicast packets per second.<br>Key: network\|port_statistics\|mcast_rx_pkts |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Ingress Broadcast Packets per second | Ingress broadcast packets per second.<br>Key: network\|port_statistics\|bcast_rx_pkts |
| Egress Dropped Packets per second | Egress dropped packets per second.<br>Key: network\|port_statistics\|dropped_tx_pkts |
| Ingress Dropped Packets per second | Ingress dropped packets per second.<br>Key: network\|port_statistics\|dropped_rx_pkts |
| Total Ingress Packets per second | Total ingress packets per second.<br>Key: network\|port_statistics\|rx_pkts |
| Total Egress Packets per second | Total egress packets per second.<br>Key: network\|port_statistics\|tx_pkts |
| Utilization | Use (KBps).<br>Key: network\|port_statistics\|utilization |
| Total Dropped Packets per second | Total dropped packets per second.<br>Key: network\|port_statistics\|dropped_pkts |
| Percentage of Dropped Packets | Percentage of dropped packets.<br>Key: network\|port_statistics\|dropped_pkts_pct |
| Max Observed Ingress Traffic (KBps) | Max observed ingress traffic (KBps).<br>Key: network\|port_statistics\|maxObserved_rx_bytes |
| Max Observed Egress Traffic (KBps) | Max observed egress traffic (KBps).<br>Key: network\|port_statistics\|maxObserved_tx_bytes |
| Max Observed Utilization (KBps) | Max observed utilization (KBps).<br>Key: network\|port_statistics\|maxObserved_utilization |

**Table 411: Summary Metrics for VMware Distributed Virtual Switches**

| Metric Name | Description |
|---|---|
| Maximum Number of Ports | Maximum number of ports.<br>Key: summary\|max_num_ports |
| Used Number of Ports | Used number of ports.<br>Key: summary\|used_num_ports |
| Number of Blocked Ports | Number of blocked ports.<br>Key: summary\|num_blocked_ports |

**Table 412: Host Metrics for VMware Distributed Virtual Switches**

| Metric Name | Description |
|---|---|
| MTU Mismatch | Maximum Transmission Unit (MTU) mismatch.<br>Key: host\|mtu_mismatch |
| Teaming Mismatch | Teaming mismatch.<br>Key: host\|teaming_mismatch |
| Unsupported MTU | Unsupported MTU.<br>Key: host\|mtu_unsupported |
| Unsupported VLANs | Unsupported VLANs.<br>Key: host\|vlans_unsupported |
| Config Out Of Sync | Config Out Of Sync. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: host\|config_outofsync |
| Number of Attached pNICs | Number of attached physical NICs.<br>Key: host\|attached_pnics |

### Distributed Virtual Port Group Metrics

The vCenter Adapter instance collects network and summary metrics for distributed virtual port groups.

Distributed Virtual Port Group metrics include badge metrics. See definitions in Badge Metrics.

**Table 413: Network Metrics for Distributed Virtual Port Groups**

| Metric Name | Description |
|---|---|
| Ingress Traffic | Ingress traffic (KBps).<br>Key: network\|port_statistics\|rx_bytes |
| Egress Traffic | Egress traffic (KBps).<br>Key: network\|port_statistics\|tx_bytes |
| Egress Unicast Packets per second | Egress unicast packets per second.<br>Key: network\|port_statistics\|ucast_tx_pkts |
| Egress Multicast Packets per second | Egress multicast packets per second.<br>Key: network\|port_statistics\|mcast_tx_pkts |
| Egress Broadcast Packets per second | Egress broadcast packets per second.<br>Key: network\|port_statistics\|bcast_tx_pkts |
| Ingress Unicast Packets per second | Ingress unicast packets per second.<br>Key: network\|port_statistics\|ucast_rx_pkts |
| Ingress Multicast Packets per second | Ingress multicast packets per second.<br>Key: network\|port_statistics\|mcast_rx_pkts |
| Ingress Broadcast Packets per second | Ingress broadcast packets per second.<br>Key: network\|port_statistics\|bcast_rx_pkts |
| Egress Dropped Packets per second | Egress dropped packets per second.<br>Key: network\|port_statistics\|dropped_tx_pkts |
| Ingress Dropped Packets per second | Ingress dropped packets per second.<br>Key: network\|port_statistics\|dropped_rx_pkts |
| Total Ingress Packets per second | Total Ingress packets per second.<br>Key: network\|port_statistics\|rx_pkts |
| Total Egress Packets per second | Total Egress packets per second.<br>Key: network\|port_statistics\|tx_pkts |
| Utilization | Utilization (KBps).<br>Key: network\|port_statistics\|utilization |
| Total Dropped Packets per second | Total dropped packets per second.<br>Key: network\|port_statistics\|dropped_pkts |
| Percentage of Dropped Packets | Percentage of dropped packets.<br>Key: network\|port_statistics\|dropped_pkts_pct |
| Max Observed Ingress Traffic (KBps) | Max observed ingress traffic (KBps).<br>Key: network\|port_statistics\|maxObserved_rx_bytes |
| Max Observed Egress Traffic (KBps) | Max observed egress traffic (KBps). |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: network\|port_statistics\|maxObserved_tx_bytes |
| Max Observed Utilization (KBps) | Max observed utilization (KBps).<br>network\|port_statistics\|maxObserved_utilization |

**Table 414: Summary Metrics for Distributed Virtual Port Groups**

| Metric Name | Description |
|---|---|
| Maximum Number of Ports | Maximum number of ports.<br>Key: summary\|max_num_ports |
| Used Number of Ports | Used number of ports.<br>Key: summary\|used_num_ports |
| Number of Blocked Ports | The number of blocked ports.<br>Key: summary\|num_blocked_ports |

### Datastore Cluster Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects profile metrics for the datastore cluster resources.

#### Profiles Metrics for Datastore Cluster Resources

Profiles metrics provide information about the profile specific capacity.

| Metric Name | Description |
|---|---|
| Profiles\|Capacity Remaining Profile (Average) | The capacity remaining in terms of fitting the average consumer.<br>Key: Profiles\|capacityRemainingProfile_<profile uuid> |
| Profiles\|Capacity Remaining Profile (<custom profile name>) | Published for custom profiles enabled from policy on Datastore Cluster Resource.<br>Key: Profiles\|capacityRemainingProfile_<profile uuid> |

#### Capacity Allocation Metrics for Datastore Cluster Resources

Capacity allocation metrics provide information about the allotment of capacity, see Capacity Analytics Generated Metrics.

### Datastore Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects capacity, device, and summary metrics for datastore objects.

Capacity metrics can be calculated for datastore objects. See Capacity Analytics Generated Metrics.

#### Capacity Metrics for Datastores

Capacity metrics provide information about datastore capacity.

| Metric Name | Description |
|---|---|
| Capacity\|Available Space (GB) | This metric shows the amount of free space that a datastore has available.<br>Use this metric to know how much storage space is unused on the datastore. Try to avoid having too little free disk space in order to accommodate unexpected storage growth on the datastore. The exact size of the datastore is based on company policy.<br><br>Key: capacity\|available_space |
| Capacity\|Provisioned (GB) | This metric shows the amount of storage that was allocated to the virtual machines.<br>Use this metric to know how much storage space is being used on the datastore.<br><br>Check the metric trend to identify spikes or abnormal growth.<br><br>Key: capacity\|provisioned |
| Capacity\|Total Capacity (GB) | This metric shows the overall size of the datastore.<br>Use this metric to know the total capacity of the datastore.<br><br>Typically the size of the datastore should not be too small. VMFS datastore size has grown over the years as virtualization matures and larger virtual machines are now onboard. Ensure that the size can handle enough virtual machines to avoid datastore sprawl. A best practice is to use 5 TB for VMFS and more for vSAN.<br><br>Key: capacity\|total_capacity |
| Capacity\|Used Space (GB) | This metric shows the amount of storage that is being used on the datastore.<br>Key: capacity\|used_space |
| Capacity\|Workload (%) | Capacity workload.<br>Key: capacity\|workload |
| Capacity\|Uncommitted Space (GB) | Uncommitted space in gigabytes.<br>Key: capacity\|uncommitted |
| Capacity\|Total Provisioned Consumer Space | Total Provisioned Consumer Space.<br>Key: capacity\|consumer_provisioned |
| Capacity\|Used Space (%) | This metric shows the amount of storage that is being used on the datastore.<br>Use this metric to know the percentage of storage space being used on the datastore.<br><br>When using this metric, verify that you have at least 20% of free storage. Less than this, and you might experience problems when a snapshot is not deleted. If you have more than 50% free storage space, you are not utilizing your storage in the best possible way.<br><br>Key: capacity\|usedSpacePct |

### Device Metrics for Datastores

Device metrics provide information about device performance.

| Metric Name | Description |
|---|---|
| Devices\|Bus Resets | This metric shows the number of bus resets in the performance interval.<br>Key: devices\|busResets_summation |
| Devices\|Commands Aborted | This metric shows the number of disk commands canceled in the performance interval.<br>Key: devices\|commandsAborted_summation |
| Devices\|Commands Issued | This metric shows the number of disk commands issued in the performance interval.<br>Key: devices\|commands_summation |
| Devices\|Read Latency (ms) | This metric shows the average time taken for a read from the perspective of a guest operating system. This metric is the sum of the Kernel Disk Read Latency and Physical Device Read Latency metrics.<br>Key: devices\|totalReadLatency_averag |
| Devices\|Kernel Disk Read Latency (ms) | Average time spent in ESX host VM Kernel per read.<br>Key: devices\|kernelReadLatency_average |
| Devices\|Kernel Write Latency (ms) | Average time spent in ESX Server VM Kernel per write.<br>Key: devices\|kernelWriteLatency_average |
| Devices\|Physical Device Read Latency (ms) | Average time taken to complete a read from the physical device.<br>Key: devices\|deviceReadLatency_average |
| Devices\|Queue Write Latency (ms) | Average time spent in the ESX Server VM Kernel queue per write.<br>Key: devices\|queueWriteLatency_average |
| Devices\|Physical Device Write Latency (ms) | Average time taken to complete a write from the physical disk.<br>Key: devices\|deviceWriteLatency_average |

### Datastore Metrics for Datastores

Datastore metrics provide information about datastore use.

| Metric Name | Description |
|---|---|
| Datastore\|Total Latency (ms) | This metric shows the adjusted read and write latency at the datastore level. Adjusted means that the latency is taking into account the number of IOs. If your IO is read-dominated, the combined value is influenced by the reads.<br>This is the average of all the VMs running in the datastore. Because it is an average, some VMs logically experience higher latency that the value shown by this metric. To see the worst latency experienced by any VM, use the Maximum VM Disk Latency metric. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Use this metric to see the performance of the datastore. It is one of two key performance indicators for a datastore, the other being the Max Read Latency. The combination of Maximum and Average gives better insight into how well the datastore is coping with the demand.<br>The number should be lower than the performance you expect.<br><br>Key: datastore\|totalLatency_average |
| Datastore\|Total Throughput (KBps) | Average use in kilobytes per second.<br>Key: datastore\|usage_average |
| Datastore\|Read Latency (ms | Average amount of time for a read operation from the datastore. Total latency = kernel latency + device latency.<br>Key: datastore\|totalReadLatency_average |
| Datastore\|Write Latency (ms) | Average amount of time for a write operation to the datastore. Total latency = kernel latency + device latency.<br>Key: datastore\|totalWriteLatency_average |
| Datastore\|Demand | Demand.<br>Key: datastore\|demand |
| Datastore\|Outstanding IO requests | OIO for datastore.<br>Key: datastore\|demand_oio |
| Datastore\|Read IOPS | This metric displays the average number of read commands issued per second during the collection interval.<br>Use this metric when the total IOPS is higher than expected. See if the metric is read or write dominated. This helps determine the cause of the high IOPS. Certain workloads such as backups, anti-virus scans, and Windows updates carry a Read/Write pattern. For example, an anti-virus scan is heavy on read since it is mostly reading the file system.<br><br>Key: datastore\|numberReadAveraged_average |
| Datastore\|Write IOPS | This metric displays the average number of write commands issued per second during the collection interval.<br>Use this metric when the total IOPS is higher than expected. Drill down to see if the metric is read or write dominated. This helps determine the cause of the high IOPS. Certain workloads such as backups, anti-virus scans, and Windows updates carry a Read/Write pattern. For example, an anti-virus scan is heavy on read since it is mostly reading the file system.<br><br>Key: datastore\|numberWriteAveraged_average |
| Datastore\|Read Throughput (KBps) | This metric displays the amount of data read in the performance interval.<br>Key: datastore\|read_average |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Datastore\|Write Throughput (KBps) | This metric displays the amount of data written to disk in the performance interval.<br>Key: datastore\|write_average |

## About Datastore Metrics for Virtual SAN

The metric named `datastore|oio|workload` is not supported on Virtual SAN datastores. This metric depends on `datastore|demand_oio`, which is supported for Virtual SAN datastores.

The metric named `datastore|demand_oio` also depends on several other metrics for Virtual SAN datastores, one of which is not supported.

- The metrics named `devices|numberReadAveraged_average` and `devices|numberWriteAveraged_average` are supported.
- The metric named `devices|totalLatency_average` is not supported.

As a result, VMware Cloud Foundation Operations does not collect the metric named `datastore|oio|workload` for Virtual SAN datastores.

## Disk Space Metrics for Datastores

Disk space metrics provide information about disk space use.

| Metric Name | Description |
|---|---|
| Diskspace\|Number of Virtual Disk | Number of virtual disks.<br>Key: diskspace\|numvmdisk |
| Diskspace\|Provisioned Space (GB) | Provisioned space in gigabytes.<br>Key: diskspace\|provisioned |
| Diskspace\|Shared Used (GB) | Shared used space in gigabytes.<br>Key: diskspace\|shared |
| Diskspace\|Snapshot Space (GB) | This metric shows the amount of space taken by snapshots on a given database.<br>Use this metric to know how much storage space is being used by virtual machine snapshots on the datastore.<br><br>Check that the snapshot is using 0 GB or minimal space. Anything over 1 GB should trigger a warning. The actual value depends on how IO intensive the virtual machines in the datastore are. Run a DT on them to detect anomaly. Clear the snapshot within 24 hours, preferably when you have finished backing up, or patching.<br><br>Key: diskspace\|snapshot |
| Diskspace\|Virtual Disk Used (GB) | Virtual disk used space in gigabytes.<br>Key: diskspace\|diskused |
| Diskspace\|Virtual machine used (GB) | Virtual machine used space in gigabytes.<br>Key: diskspace\|used |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Diskspace\|Total disk space used | Total disk space used on all datastores visible to this object.<br>Key: diskspace\|total_usage |
| Diskspace\|Total disk space | Total disk space on all datastores visible to this object.<br>Key: diskspace\|total_capacity |
| Diskspace\|Total used (GB) | Total used space in gigabytes.<br>Key: diskspace\|disktotal |
| Diskspace\|Swap File Space (GB) | Swap file space in gigabytes.<br>Key: diskspace\|swap |
| Diskspace\|Other VM Space (GB) | Other virtual machine space in gigabytes.<br>Key: diskspace\|otherused |
| Diskspace\|Freespace (GB) | Unused space available on datastore.<br>Key: diskspace\|freespace |
| Diskspace\|Capacity (GB) | Total capacity of datastore in gigabytes.<br>Key: diskspace\|capacity |
| Diskspace\|Overhead | Amount of disk space that is overhead.<br>Key: diskspace\|overhead |

**Summary Metrics for Datastores**

Summary metrics provide information about overall performance.

| Metric Name | Description |
|---|---|
| Summary\|Number of Hosts | This metric shows the number of hosts that the datastore is connected to.<br>Use this metric to know how many clusters the datastore is attached to.<br><br>The number should not be too high, as a datastore should not be mounted by every host. The datastore and cluster should be paired to keep operations simple.<br><br>Key: summary\|total_number_hosts |
| Summary\|Total Number of VMs | This metric shows the number of virtual machines which save their VMDK files on the datastore. If a VM has four VMDKs stored in four datastores, the VM is counted on each datastore.<br>Use this metric to know how many VMs have at least one VMDK on a specific datastore.<br><br>The number of VMs should be within your Concentration Risk policy.<br><br>You should also expect the datastore to be well used. If only a few VMs are using the datastore, this is not considered a good use. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: summary\|total_number_vms |
| Summary\|Maximum Number of VMs | Maximum number of virtual machines.<br>Key: summary\|max_number_vms |
| Summary\|Workload Indicator | Workload indicator.<br>Key: summary\|workload_indicator |
| Summary\|Number of Clusters | This metric shows the number of clusters that the datastore is connected to.<br>Key: summary\|total_number_clusters |
| Summary\|Number of VM Templates | Number of VM Templates.<br>Key: Summary\|Number of VM Templates |

**Template Metrics for Datastores**

| Metric Name | Description |
|---|---|
| Template\|Virtual Machine used | Space used by virtual machine files.<br>Key: template\|used |
| Template\|Access Time | Last access time.<br>Key: template\|accessTime |

**Cost Metrics for Datastores**

Cost metrics provides information about the cost.

| Metric Name | Description |
|---|---|
| Monthly Disk Space Base Rate | Disk space base rate for datastore displays the cost of 1 GB storage.<br>Key: cost\|storageRate |
| Monthly Total Cost | Monthly total cost, computed by multiplying datastore capacity with monthly storage rate.<br>Key: cost\|totalCost |
| Cost\|Allocation\|Disk Space Base Rate (Currency) | Monthly storage rate for datastore displays the cost of 1 GB storage when the overcommit ratio is set in policy.<br>cost\|storageRate |
| Cost\|Allocation\|Monthly Datastore Allocated Cost(Currency/Month) | Monthly allocated cost as compared to the total cost of the datastore |
| Cost\|Allocation\|Monthly Datastore Unallocated Cost(Currency/Month) | Monthly unallocated cost as compared to the total cost of the datastore. |

**Reclaimable Metrics**

Reclaimable metrics provide information about reclaimable resources.

| Metric Name | Description |
|---|---|
| Reclaimable\|Orphaned Disks\|Disk Space (GB) | Summary of storage used by all orphaned VMDKs on the datastore.<br>Key: reclaimable\|orphaned_disk\|diskspace |
| Reclaimable\|Orphaned Disks\|Potential Savings (Currency) | Potential saving after reclamation of storage by removing orphaned VMDKs from the datastore.<br>Key: reclaimable\|orphaned_disk\|cost |

**Disabled Instanced Metrics**

The instance metrics created for the following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that these metrics collect data by default but all the instanced metrics created for these metrics, do not collect data by default.

| Metric Name |
|---|
| Devices\|Kernel Latency (ms) |
| Devices\|Number of Running Hosts |
| Devices\|Number of Running VMs |
| Devices\|Physical Device Latency (ms) |
| Devices\|Queue Latency (ms) |
| Devices\|Queue Read Latency (ms) |
| Devices\|Read IOPS |
| Devices\|Read Latency (ms) |
| Devices\|Read Requests |
| Devices\|Read Throughput (KBps) |
| Devices\|Total IOPS |
| Devices\|Total Latency (ms) |
| Devices\|Total Throughput (KBps) |
| Devices\|Write IOPS |
| Devices\|Write Latency (ms) |
| Devices\|Write Requests |
| Devices\|Write Throughput (KBps) |

**Disabled Metrics**

The following metrics are disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

You can enable these metrics in the Policy workspace. For more information, in VMware Docs search for Collect Metrics and Properties Details.

You can enable these metrics in the Policy workspace. For more information, see Metrics and Properties Details.

| Metric Name | Key |
|---|---|
| Capacity\|Data Store Capacity Contention (%) | capacity\|contention |
| Datastore I/O\|Demand Indicator | datastore\|demand_indicator |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Key |
|---|---|
| Datastore I/O\|Max Observed Number of Outstanding IO Operations | datastore\|maxObserved_OIO |
| Datastore I/O\|Max Observed Read Latency (msec) | datastore\|maxObserved_Read |
| Datastore I/O\|Max Observed Read Latency (msec) | datastore\|maxObserved_ReadLatency |
| Datastore I/O\|Max Observed | datastore\|maxObserved_NumberRead |
| Datastore I/O\|Max Observed Write Latency (msec) | datastore\|maxObserved_Write |
| Datastore I/O\|Max Observed Write Latency (msec) | datastore\|maxObserved_WriteLatency |
| Datastore I/O\|Max Observed Writes per second | datastore\|maxObserved_NumberWrite |
| Datastore\|Demand Indicator | Demand Indicator.<br>Key: datastore\|demand_indicator |
| Diskspace\|Not Shared (GB) | Unshared space in gigabytes.<br>Key: diskspace\|notshared |

### Cluster Compute Metrics for Allocation Model

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, disk space, CPU use, disk, memory, network, power, and summary metrics for cluster compute resources.

#### Cost Metrics for Cluster Compute Resources

Cost metrics provide information about the cost.

| Metric Name | Description |
|---|---|
| Cluster CPU Base Rate | Base rate for Cluster CPU calculated by dividing the monthly total cluster CPU cost by cluster CPU over-commit ratio.<br>Key:Cost\|Allocation\|ClusterCPUBaseRate |
| Cluster Memory Base Rate | Cluster memory base rate calculated by dividing the monthly total cluster memory cost b cost by cluster memory over-commit ratio.<br>Key: Cost\|Allocation\|ClusterMemoryBaseRate |
| Monthly Cluster Allocated Cost | Sum of of monthly cluster CPU, Memory, and Storage costs<br>Key: Cost\|Allocation\|MonthlyClusterAllocatedCost |
| Monthly Cluster Unallocated Cost | Monthly cluster unallocated cost calculated by subtracting the monthly cluster allocated cost from the monthly cluster total cost.<br>Key: Cost\|Allocation\| MonthlyClusterUnallocatedCost |
| Monthly Storage Rate | Datastore base rate is calculated by dividing Storage base rate based on utilization by over commit ratio.<br>Key:Cost\|Allocation\|Monthly Storage Rate |

### Virtual Machine Metrics for Allocation Model

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, disk space, CPU use, disk, memory, network, power, and summary metrics for virtual machine resources.

#### Cost Metrics for Virtual Machines

Cost metrics provide information about the cost.

| Metric Name | Description |
|---|---|
| MTD VM CPU Cost | Month to date virtual machine CPU cost.<br>Key: Cost\|Allocation\|MTD VM CPU Cost |
| MTD VM Memory Cost | Month to date virtual machine memory cost.<br>Key: Cost\|Allocation\|MTD VM Memory Cost |
| MTD VM Storage Cost | Month to date storage cost of the virtual machine.<br>Key: Cost\|Allocation\|MTD VM Storage Cost |
| MTD VM Total Cost | Addition of CPU ,Memory ,Storage, and Direct cost.<br>Key: Cost\|Allocation\|MTD VM Total Cost |

### Metrics for Namespace

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for Namspace through the vCenter adapter and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

**Table 415: Metrics for Namespace**

| Metric Key | Localized Name | Description |
|---|---|---|
| cpu\|usagemhz_average | CPU\|Usage | Average CPU usage in MHZ. |
| cpu\|demandmhz | CPU\|Demand | Demand(MHz). |
| cpu\|capacity_contentionPct | CPU\|Contention | Percent of time descendant virtual machines are unable to run because they are contending for access to the physical CPU(s). |
| cpu\|effective_limit | CPU\|Effective limit | CPU Effective limit. |
| cpu\|reservation_used | CPU\|Reservation Used | CPU Reservation Used. |
| cpu\|estimated_entitlement | CPU\|Estimated entitlement | CPU Estimated entitlement. |
| cpu\|dynamic_entitlement | CPU\|Dynamic entitlement | CPU Dynamic Entitlement. |
| cpu\|capacity_contention | CPU\|Overall CPU Contention | Overall CPU Contention (ms). |
| cpu\|<br>capacity_demandEntitlementPct | CPU\|Capacity Demand Entitlement | CPU Capacity Demand Entitlement Percentage. |
| mem\|usage_average | Memory\|Usage | Memory currently in use as a percentage of total available memory. |
| mem\|guest_provisioned | Memory\|Total Capacity | Total Capacity. |
| mem\|active_average | Memory\|Guest Active | Amount of memory that is actively used. |
| mem\|granted_average | Memory\|Granted | Amount of memory available for use. |
| mem\|shared_average | Memory\|Shared | Amount of shared memory. |
| mem\|overhead_average | Memory\|VM Overhead | Memory overhead reported by host. |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Localized Name | Description |
|---|---|---|
| mem\|consumed_average | Memory\|Consumed | Amount of host memory consumed by the virtual machine for guest memory. |
| mem\|host_contentionPct | Memory\|Contention | Machine Contention Percentage. |
| mem\|guest_usage | Memory\|Guest Usage | Guest Memory Entitlement. |
| mem\|guest_demand | Memory\|Guest Demand | Guest Memory Entitlement. |
| mem\|reservation_used | Memory\|Reservation Used | Memory Reservation Used. |
| mem\|effective_limit | Memory\|Effective limit | Memory Effective limit. |
| mem\|swapinRate_average | Memory\|Swap In Rate | Rate at which memory is swapped from disk into active memory during the collection interval. This can impact performance. |
| mem\|swapoutRate_average | Memory\|Swap Out Rate | Rate at which memory is being swapped from active memory to disk during the current interval. |
| mem\|vmmemctl_average | Memory\|Balloon | Amount of memory currently used by the virtual machine memory control. |
| mem\|zero_average | Memory\|Zero | Amount of memory that is all 0. |
| mem\|swapped_average | Memory\|Swapped | Amount of unreserved memory. |
| mem\|zipped_latest | Memory\|Zipped | N/A |
| mem\|compressionRate_average | Memory\|Compression Rate | N/A |
| mem\|decompressionRate_average | Memory\|Decompression Rate | N/A |
| mem\|swapin_average | Memory\|Swap In | Amount of memory swapped in. |
| mem\|swapout_average | Memory\|Swap Out | Amount of memory swapped out. |
| mem\|swapused_average | Memory\|Swap Used | Amount of memory used for swap space. |
| mem\|host_contention | Memory\|Contention | Machine Contention. |
| mem\|dynamic_entitlement | Memory\|Dynamic Entitlement | Memory Dynamic Entitlement. |
| diskspace\|total_usage | Disk Space\|Utilization | Storage space utilized on connected vSphere Datastores. |
| summary\|configStatus | Summary\|Config Status | Workload Management Configuration Status. |
| summary\|total_number_pods | Summary\|Number of Pods | Number of Pods. |
| summary\| numberKubernetesClusters | Summary\|Number of Kubernetes clusters | Number of Kubernetes clusters. |
| summary\|number_running_vms | Summary\|Number of Running VMs | Number of Running VMs. |
| summary\|total_number_vms | Summary\|Total Number of VMs | Total Number of VMs. |
| summary\|iowait | Summary\|IO Wait | IO Wait. |

### Metrics for Tanzu Kubernetes cluster

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for Tanzu Kubernetes cluster through the vCenter adapter and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

**Table 416: Metrics for Tanzu Kubernetes clusters**

| Metric Key | Localized Name | Description |
|---|---|---|
| cpu\|usagemhz_average | CPU\|Usage | Average CPU usage in MHZ |
| cpu\|demandmhz | CPU\|Demand | Demand(MHz) |
| cpu\|capacity_contentionPct | CPU\|Contention | Percent of time descendant virtual machines are unable to run because they are contending for access to the physical CPU(s). |
| cpu\|effective_limit | CPU\|Effective limit | CPU Effective limit |
| cpu\|reservation_used | CPU\|Reservation Used | CPU Reservation Used |
| cpu\|estimated_entitlement | CPU\|Estimated entitlement | CPU Estimated entitlement |
| cpu\|dynamic_entitlement | CPU\|Dynamic entitlement | CPU Dynamic Entitlement |
| cpu\|capacity_contention | CPU\|Overall CPU Contention | Overall CPU Contention (ms) |
| cpu\|capacity_demandEntitlementPct | CPU\|Capacity Demand Entitlement | CPU Capacity Demand Entitlement Percentage |
| mem\|usage_average | Memory\|Usage | Memory currently in use as a percentage of total available memory |
| mem\|guest_provisioned | Memory\|Total Capacity | Total Capacity |
| mem\|active_average | Memory\|Guest Active | Amount of memory that is actively used |
| mem\|granted_average | Memory\|Granted | Amount of memory available for use |
| mem\|shared_average | Memory\|Shared | Amount of shared memory |
| mem\|overhead_average | Memory\|VM Overhead | Memory overhead reported by host |
| mem\|consumed_average | Memory\|Consumed | Amount of host memory consumed by the virtual machine for guest memory |
| mem\|host_contentionPct | Memory\|Contention | Machine Contention Percentage |
| mem\|guest_usage | Memory\|Guest Usage | Guest Memory Entitlement |
| mem\|guest_demand | Memory\|Guest Demand | Guest Memory Entitlement |
| mem\|reservation_used | Memory\|Reservation Used | Memory Reservation Used |
| mem\|effective_limit | Memory\|Effective limit | Memory Effective limit |
| mem\|swapinRate_average | Memory\|Swap In Rate | Rate at which memory is swapped from disk into active memory during the collection interval. This can impact performance. |
| mem\|swapoutRate_average | Memory\|Swap Out Rate | Rate at which memory is being swapped from active memory to disk during the current interval |
| mem\|vmmemctl_average | Memory\|Balloon | Amount of memory currently used by the virtual machine memory control |
| mem\|zero_average | Memory\|Zero | Amount of memory that is all 0 |
| mem\|swapped_average | Memory\|Swapped | Amount of unreserved memory |
| mem\|zipped_latest | Memory\|Zipped | N/A |
| mem\|compressionRate_average | Memory\|Compression Rate | N/A |
| mem\|decompressionRate_average | Memory\|Decompression Rate | N/A |
| mem\|swapin_average | Memory\|Swap In | Amount of memory swapped in |
| mem\|swapout_average | Memory\|Swap Out | Amount of memory swapped out |
| mem\|swapused_average | Memory\|Swap Used | Amount of memory used for swap space |
| mem\|host_contention | Memory\|Contention | Machine Contention |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Localized Name | Description |
|---|---|---|
| mem\|dynamic_entitlement | Memory\|Dynamic Entitlement | Memory Dynamic Entitlement |
| summary\|number_running_vms | Summary\|Number of Running VMs | Number of Running VMs |
| summary\|total_number_vms | Summary\|Total Number of VMs | Total Number of VMs |
| summary\|iowait | Summary\|IO Wait | IO Wait |

### Metrics for vSphere Pods

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for vSphere Pods through the vCenter adapter and uses formulas to derive statistics from those metrics. You can use metrics to troubleshoot problems in your environment.

**Table 417: Metrics for vSphere Pods**

| Metric Key | Metric Name | Description |
|---|---|---|
| config\|hardware\|num_Cpu | Configuration\|Hardware\|Number of CPUs | Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU. |
| config\|hardware\|disk_Space | Configuration\|Hardware\|Disk Space | Disk space metrics |
| config\|hardware\|thin_Enabled | Configuration\|Hardware\|Thin Provisioned Disk | Thin Provisioned Disk |
| config\|cpuAllocation\|slotSize | Configuration\|CPU Resource Allocation\|HA Slot Size | vSphere HA Slot Size for CPU |
| config\|memoryAllocation\|slotSize | Configuration\|Memory Resource Allocation\|HA Slot Size | vSphere HA Slot Size for Memory |
| cpu\|usage_average | CPU\|Usage | CPU Usage divided by VM CPU Configuration in MHz |
| cpu\|usagemhz_average | CPU\|Usage | Amount of actively used virtual CPU. This is the host's view of the CPU usage, not the guest operating system view. |
| cpu\|usagemhz_average_mtd | CPU\|Usage average MTD | Month to date average CPU usage in MHZ |
| cpu\|readyPct | CPU\|Ready | Percentage of CPU the VM is ready to run, but unable due to ESXi has no ready physical core to run it. High Ready value impacts VM performance |
| cpu\|capacity_contentionPct | CPU\|Contention | Percentage of time VM is not getting the CPU resource it demanded. Impacted by Ready, Co-Stop, Hyper Threading and Power Management |
| cpu\|corecount_provisioned | CPU\|Provisioned vCPU(s) | Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU. |
| cpu\|vm_capacity_provisioned | CPU\|Total Capacity | Configured Capacity in MHz, based on nominal (static) frequency of the CPU |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| cpu\|demandmhz | CPU\|Demand | The amount of CPU resources virtual machine would use if there were no CPU contention or CPU limit. |
| cpu\|demandPct | CPU\|Demand (%) | The percentage of CPU resources virtual machine would use if there were no CPU contention or CPU limit. |
| cpu\|reservation_used | CPU\|Reservation Used | CPU Reserved for the VM. It's guaranteed to be available when the VM demands it. |
| cpu\|effective_limit | CPU\|Effective limit | Limit placed on the VM by vSphere. Avoid using limit as it impacts VM performance |
| cpu\|iowaitPct | CPU\|IO Wait | Percentage of time VM CPU is waiting for IO. Formula is Wait - Idle - Swap Wait. High value indicates slow storage subsystem |
| cpu\|swapwaitPct | CPU\|Swap wait | Percentage of time CPU is waiting on data swap-in. Mapped to vCenter CPU Swap wait |
| cpu\|costopPct | CPU\|Co-stop (%) | Percentage of time the VM is ready to run, but is unable to due to co-scheduling constraints. VM with less vCPU have lower co-stop value. |
| cpu\|system_summation | CPU\|System | CPU time spent on system processes |
| cpu\|wait_summation | CPU\|Wait | Total CPU time spent in wait state |
| cpu\|ready_summation | CPU\|Ready | CPU time spent on ready state |
| cpu\|used_summation | CPU\|Used | CPU time that is used |
| cpu\|iowait | CPU\|IO Wait | IO Wait |
| cpu\|wait | CPU\|Total Wait | CPU time spent on idle state |
| cpu\|<br>capacity_demandEntitlementPct | CPU\|Capacity Demand Entitlement | CPU Capacity Demand Entitlement Percentage |
| cpu\|host_demand_for_aggregation | CPU\|Host Demand For Aggregation | Host demand for aggregation |
| cpu\|dynamic_entitlement | CPU\|Dynamic entitlement | CPU Dynamic entitlement |
| cpu\|capacity_contention | CPU\|Overall CPU Contention | Overall CPU Contention (ms) |
| cpu\|estimated_entitlement | CPU\|Estimated entitlement | CPU Estimated entitlement |
| cpu\|idlePct | CPU\|Idle | % CPU time that is idle |
| cpu\|waitPct | CPU\|Wait | % Total CPU time spent in wait state |
| cpu\|systemSummationPct | CPU\|System | % CPU time spent on system processes |
| cpu\|demandOverLimit | CPU\|Demand Over Limit | Amount of CPU Demand that is over the configured CPU Limit |
| cpu\|demandOverCapacity | CPU\|Demand Over Capacity | Amount of CPU Demand that is over the configured CPU Capacity |
| cpu\|perCpuCoStopPct | CPU\|Normalized Co-stop | Percentage of co-stop time, normalized across all vCPUs |
| cpu\|swapwait_summation | CPU\|Swap Wait | Amount of time waiting on swap. |
| cpu\|costop_summation | CPU\|Co-stop | Time the VM is ready to run, but is unable to due to co-scheduling constraints. |
| cpu\|idle_summation | CPU\|Idle | CPU time that is idle. |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| cpu\|latency_average | CPU\|Latency | Percentage of time the VM is unable to run because it is contending for access to the physical CPUs. |
| cpu\|maxlimited_summation | CPU\|Max Limited | Time the VM is ready to run, but is not run due to maxing out its CPU limit setting. |
| cpu\|overlap_summation | CPU\|Overlap | Time the VM was interrupted to perform system services on behalf of that VM or other VMs. |
| cpu\|run_summation | CPU\|Run | Time the VM is scheduled to run. |
| cpu\|entitlement_latest | CPU\|Entitlement Latest | Entitlement Latest. |
| cpu\|demandEntitlementRatio_latest | CPU\|Demand-to-entitlement Ratio | CPU resource entitlement to CPU demand ratio (in percents) |
| cpu\|readiness_average | CPU\|Readiness | Percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU. |
| rescpu\|actav1_latest | CPU Utilization for Resources\| CPU Active (1 min. average) | The average active time for the CPU over the past minute |
| rescpu\| actav5_latestswapinRate_average | CPU Utilization for Resources\| CPU Active (5 min. average) | The average active time for the CPU over the past five minutes. |
| rescpu\|actav5_latest | CPU Utilization for Resources\| CPU Active (5 min. average) | The average active time for the CPU over the past five minutes |
| rescpu\|actav15_latest | CPU Utilization for Resources\| CPU Active (15 min. average) | The average active time for the CPU over the past fifteen minutes |
| rescpu\|actpk1_latest | CPU Utilization for Resources\| CPU Active (1 min. peak) | The peak active time for the CPU over the past minute |
| rescpu\|actpk5_latest | CPU Utilization for Resources\| CPU Active (5 min. peak) | The peak active time for the CPU over the past five minutes |
| rescpu\|actpk15_latest | CPU Utilization for Resources\| CPU Active (15 min. peak) | The peak active time for the CPU over the past fifteen minutes |
| rescpu\|runav1_latest | CPU Utilization for Resources\| CPU Running (1 min. average) | The average runtime for the CPU over the past minute |
| rescpu\|runav5_latest | CPU Utilization for Resources\| CPU Running (5 min. average) | The average runtime for the CPU over the past five minutes |
| rescpu\|runav15_latest | CPU Utilization for Resources\| CPU Running (15 min. average) | The average runtime for the CPU over the past fifteen minutes |
| rescpu\|runpk1_latest | CPU Utilization for Resources\| CPU Running (1 min. peak) | The peak active time for the CPU over the past minute |
| rescpu\|runpk5_latest | CPU Utilization for Resources\| CPU Running (5 min. peak) | The peak active time for the CPU over the past five minutes |
| rescpu\|runpk15_latest | CPU Utilization for Resources\| CPU Running (15 min. peak) | The peak active time for the CPU over the past fifteen minutes |
| rescpu\|maxLimited1_latest | CPU Utilization for Resources\| CPU Throttled (1 min. average) | The scheduling limit over the past minute |
| rescpu\|maxLimited5_latest | CPU Utilization for Resources\| CPU Throttled (5 min. average) | The scheduling limit over the past five minutes |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| rescpu\|maxLimited15_latest | CPU Utilization for Resources\| CPU Throttled (15 min. average) | The scheduling limit over the past fifteen minutes |
| rescpu\|sampleCount_latest | CPU Utilization for Resources\| Group CPU Sample Count | The sample CPU count |
| rescpu\|samplePeriod_latest | CPU Utilization for Resources\| Group CPU Sample Period | The sample period |
| mem\|usage_average | Memory\|Usage | Memory currently in use as a percentage of total available memory |
| mem\|balloonPct | Memory\|Balloon | Percentage of guest physical memory that is currently claimed from the virtual machine through ballooning. This is the percentage of guest physical memory that has been allocated and pinned by the balloon driver. Balloon does not necessarily mean the VM performance is affected. |
| mem\|swapped_average | Memory\|Swapped | Amount of unreserved memory |
| mem\|consumed_average | Memory\|Consumed | Amount of ESXi Host memory mapped/ consumed by the virtual machine for guest memory |
| mem\|consumed_average_mtd | Memory\|Consumed average MTD | average MTD Amount of host memory consumed by the virtual machine for guest memory |
| mem\|consumedPct | Memory\|Consumed (%) | Amount of host memory consumed by the virtual machine for guest memory. Consumed memory does not include overhead memory. It includes shared memory and memory that might be reserved, but not actually used. |
| mem\|overhead_average | Memory\|Overhead | Amount of overhead memory used by ESXi to run the Virtual Machine. |
| mem\|host_contentionPct | Memory\|Contention | Percentage of time the VM has contended for memory. |
| mem\|guest_provisioned | Memory\|Total Capacity | Memory resources allocated to the Virtual Machine |
| mem\|guest_usage | Memory\|Guest Usage | Guest Memory Entitlement |
| mem\|guest_demand | Memory\|Guest Demand | Guest Memory Entitlement |
| mem\|host_demand | Memory\|Host Demand | Memory Demand in KB |
| mem\|reservation_used | Memory\|Reservation Used | Memory Reservation Used |
| mem\|effective_limit | Memory\|Effective limit | Memory Effective limit |
| mem\|vmMemoryDemand | Memory\|Utilization | Amount of memory utilized by the Virtual Machine. Reflects the guest OS memory required (for certain vSphere and VMTools versions) or Virtual Machine consumption |
| mem\|nonzero_active | Memory\|Non Zero Active | Non Zero Active Memory |
| mem\|swapinRate_average | Memory\|Swap In Rate | Rate at which memory is swapped from disk into active memory during the collection interval. This can impact performance. |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| mem\|swapoutRate_average | Memory\|Swap Out Rate | Rate at which memory is being swapped from active memory to disk during the current interval. |
| mem\|compressed_average | Memory\|Compressed | Percentage of total memory that has been compressed by vSphere. If and only if the page is accessed by the Guest OS, will performance be affected. |
| mem\|overheadMax_average | Memory\|Overhead Max | N/A |
| mem\|vmmemctl_average | Memory\|Balloon | Amount of memory currently used by the virtual machine memory control |
| mem\|active_average | Memory\|Guest Active | Amount of memory that is actively used |
| mem\|granted_average | Memory\|Granted | Amount of memory available for use |
| mem\|shared_average | Memory\|Shared | Amount of shared memory |
| mem\|zero_average | Memory\|Zero | Amount of memory that is all 0 |
| mem\|swaptarget_average | Memory\|Swap Target | Amount of memory that can be swapped |
| mem\|swapin_average | Memory\|Swap In | Amount of memory swapped in |
| mem\|swapout_average | Memory\|Swap Out | Amount of memory swapped out |
| mem\|vmmemctltarget_average | Memory\|Balloon Target | Amount of memory that can be used by the virtual machine memory control |
| mem\|host_dynamic_entitlement | Memory\|Host Dynamic Entitlement | Mem Machine Dynamic Entitlement |
| mem\|host_active | Memory\|Host Active | Machine Active |
| mem\|host_usage | Memory\|Host Usage | Machine Usage |
| mem\|host_contention | Memory\|Contention | Machine Contention |
| mem\|guest_activePct | Memory\|Guest Active Memory | Guest active memory as percentage of configured |
| mem\|guest_dynamic_entitlement | Memory\|Guest Dynamic Entitlement | Guest Memory Dynamic Entitlement |
| mem\|host_demand_reservation | Memory\|Host Demand with Reservation | Memory Demand with Reservation considered in KB |
| mem\|host_nonpageable_estimate | Memory\|Guest Non Pageable Memory | Guest Non Pageable Memory Estimates |
| mem\|guest_nonpageable_estimate | Memory\|Host Non Pageable Memory | Guest Non Pageable Memory Estimates |
| mem\|estimated_entitlement | Memory\|Estimated entitlement | Memory Estimated entitlement |
| mem\| host_demand_for_aggregation | Memory\|Host Demand For Aggregation | Host demand for aggregation |
| mem\|demandOverLimit | Memory\|Demand Over Limit | Amount of Memory Demand that is over the configured Memory Limit |
| mem\|demandOverCapacity | Memory\|Demand Over Capacity | Amount of Memory Demand that is over the configured Memory Capacity |
| mem\|activewrite_average | Memory\|Active Write | N/A |
| mem\|compressionRate_average | Memory\|Compression Rate | N/A |
| mem\|decompressionRate_average | Memory\|Decompression Rate | N/A |
| mem\|zipSaved_latest | Memory\|Zip Saved | N/A |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| mem\|zipped_latest | Memory\|Zipped | N/A |
| mem\|entitlement_average | Memory\|Entitlement | Amount of host physical memory the VM is entitled to, as determined by the ESX schedule. |
| mem\|latency_average | Memory\|Latency | Percentage of time the VM is waiting to access swapped or compressed memory. |
| mem\|capacity.contention_average | Memory\|Capacity Contention | Capacity Contention. |
| mem\|llSwapInRate_average | Memory\|Swap In Rate from Host Cache | Rate at which memory is being swapped from host cache into active memory. |
| mem\|llSwapOutRate_average | Memory\|Swap Out Rate to Host Cache | Rate at which memory is being swapped to host cache from active memory. |
| mem\|llSwapUsed_average | Memory\|Swap Space Used in Host Cache | Space used for caching swapped pages in the host cache. |
| mem\|overheadTouched_average | Memory\|Overhead Touched | Actively touched overhead memory (KB) reserved for use as the virtualization overhead for the VM. |
| net\|usage_average | Network\|Usage Rate | The sum of the data transmitted and received for all the NIC instances of the host or virtual machine |
| net\|transmitted_average | Network\|Data Transmit Rate | Average amount of data transmitted per second |
| net\|received_average | Network\|Data Receive Rate | Average amount of data received per second |
| net\|droppedTx_summation | Network\|Transmitted Packets Dropped | Number of outgoing packets dropped in the performance interval. Investigate if the number is not 0 |
| net\|droppedPct | Network\|Packets Dropped (%) | Percentage of packets dropped |
| net\|dropped | Network\|Packets Dropped | Number of packets dropped in the performance interval |
| net\|broadcastTx_summation | Network\|Broadcast Packets Transmitted | Total number of broadcast packets transmitted. Investigate further if this number is high |
| net\|multicastTx_summation | Network\|Multicast Packets Transmitted | Number of multicast packets transmitted. Investigate further if this number is high |
| net\|idle | Networkldle | N/A |
| net\|usage_capacity | Network\|I/O Usage Capacity | I/O Usage Capacity |
| net\|maxObserved_KBps | Network\|Max Observed Throughput | Max observed rate of network throughput |
| net\|maxObserved_Tx_KBps | Network\|Max Observed Transmitted Throughput | Max observed transmitted rate of network throughput |
| net\|maxObserved_Rx_KBps | Network\|Max Observed Received Throughput | Max observed received rate of network throughput |
| net\|packetsRx_summation | Network\|Packets Received | Number of packets received in the performance interval |
| net\|packetsTx_summation | Network\|Packets Transmitted | Number of packets transmitted in the performance interval |
| net\|demand | Network\|Demand | N/A |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| net\|packetsRxPerSec | Network\|Packets Received per second | Number of packets received in the performance interval |
| net\|packetsTxPerSec | Network\|Packets Transmitted per second | Number of packets transmitted in the performance interval |
| net\|packetsPerSec | Network\|Packets per second | Number of packets transmitted and received per second |
| net\|droppedRx_summation | Network\|Received Packets Dropped | Number of received packets dropped in the performance interval |
| net\|broadcastRx_summation | Network\|Broadcast Packets Received | Number of broadcast packets received during the sampling interval |
| net\|multicastRx_summation | Network\|Multicast Packets Received | Number of multicast packets received |
| net\|bytesRx_average | Network\|bytesRx | Average amount of data received per second |
| net\|bytesTx_average | Network\|bytesTx | Average amount of data transmitted per second |
| net\|host_transmitted_average | Network\|VM to Host Data Transmit Rate | Average amount of data transmitted per second between VM and host |
| net\|host_received_average | Network\|VM to Host Data Receive Rate | Average amount of data received per second between VM and host |
| net\|host_usage_average | Network\|VM to Host Usage Rate | The sum of the data transmitted and received for all the NIC instances between VM and host |
| net\|host_maxObserved_Tx_KBps | Network\|VM to Host Max Observed Transmitted Throughput | Max observed transmitted rate of network throughput between VM and host |
| net\|host_maxObserved_Rx_KBps | Network\|VM to Host Max Observed Received Throughput | Max observed received rate of network throughput between VM and host |
| net\|host_maxObserved_KBps | Network\|VM to Host Max Observed Throughput | Max observed rate of network throughput between VM and host |
| net\|transmit_demand_average | Network\|Data Transmit Demand Rate | Data Transmit Demand Rate |
| net\|receive_demand_average | Network\|Data Receive Demand Rate | Data Receive Demand Rate |
| disk\|usage_average | Physical Disk\|Total Throughput | Amount of data read from/written to storage in a second. This is averaged over the reporting period |
| disk\|read_average | Physical Disk\|Read Throughput | Amount of data read from storage in a second. This is averaged over the reporting period |
| disk\|write_average | Physical Disk\|Write Throughput | Amount of data written to storage in a second. This is averaged over the reporting period |
| disk\|usage_capacity | Physical Disk\|I/O Usage Capacity | I/O Usage Capacity |
| disk\|busResets_summation | Physical Disk\|Bus Resets | The number of bus resets in the performance interval |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| disk\|commandsAborted_summation | Physical Disk\|Commands Aborted | The number of disk commands stopped in the performance interval |
| disk\|diskoio | Physical Disk\|Number of Outstanding IO Operations | Number of Outstanding IO Operations |
| disk\|diskqueued | Physical Disk\|Queued Operations | Queued Operations |
| disk\|diskdemand | Physical Disk\|Demand | Demand |
| disk\|sum_queued_oio | Physical Disk\|Total Queued Outstanding operations | Sum of Queued Operation and Outstanding Operations. |
| disk\|max_observed | Physical Disk\|Max Observed OIO | Max Observed IO for a disk. |
| disk\| numberReadAveraged_average | Physical Disk\|Read IOPS | Number of read operations per second. This is averaged over the reporting period. |
| disk\| numberWriteAveraged_average | Physical Disk\|Write IOPS | Number of write operations per second. This is averaged over the reporting period. |
| disk\|maxTotalLatency_latest | Physical Disk\|Highest Latency | Highest Latency. |
| disk\| scsiReservationConflicts_summation | Physical Disk\|SCSI Reservation Conflicts | SCSI Reservation Conflicts. |
| disk\|totalReadLatency_average | Physical Disk\|Read Latency | Average amount of time for a read operation by the storage adapter. |
| disk\|totalWriteLatency_average | Physical Disk\|Write Latency | Average amount of time for a write operation by the storage adapter. |
| disk\|totalLatency_average | Physical Disk\|Total Latency | Total Latency. |
| sys\|poweredOn | System\|Powered ON | 1 if the VM is connected (available for management) and powered on, otherwise 0. |
| sys\|osUptime_latest | System\|OS Uptime | Total time elapsed, in seconds, since last operating system boot-up |
| sys\|uptime_latest | System\|Uptime | Number of seconds since system startup |
| sys\|heartbeat_summation | System\|Heartbeat | Number of heart beats from the virtual machine in the defined interval |
| sys\|vmotionEnabled | System\|vMotion Enabled | 1 if vMotion enabled, 0 if not enabled |
| sys\|productString | System\|Product String | VMware product string |
| sys\|heartbeat_latest | System\|Heartbeat Latest | Number of heartbeats issued per virtual machine during the interval |
| summary\|running | Summary\|Running | Running |
| summary\|desktop_status | Summary\|Desktop Status | Horizon View Desktop Status |
| summary\|poweredOff | Summary\|Reclaimable Powered Off | Powered Off = 1. Not powered off = 0 |
| summary\|idle | Summary\|Reclaimable Idle | Idle = 1. Not idle = 0 |
| summary\|oversized | Summary\|Is Oversized | Oversized = 1. Not oversized = 0 |
| summary\|undersized | Summary\|Is Undersized | Is Undersized |
| summary\|snapshotSpace | Summary\|Reclaimable Snapshot Space | Reclaimable Snapshot Space |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| summary\|oversized\|vcpus | Summary\|Oversized\|Virtual CPUs | Virtual CPUs |
| summary\|oversized\|memory | Summary\|Oversized\|Memory | Memory |
| summary\|undersized\|vcpus | Summary\|Undersized\|Virtual CPUs | Virtual CPUs |
| summary\|undersized\|memory | Summary\|Undersized\|Memory | Memory |
| summary\|metering\|value | Summary\|Metering\|Total price | Total price of the resource(Sum of all price components) |
| summary\|metering\|storage | Summary\|Metering\|Storage price | Price of Storage related components of the resource |
| summary\|metering\|memory | Summary\|Metering\|Memory price | Price of Memory related components of the resource |
| summary\|metering\|cpu | Summary\|Metering\|CPU price | Price of CPU related components of the resource |
| summary\|metering\|additional | Summary\|Metering\|Additional price | Price of additional components of the resource |
| summary\|metering\|partialPrice | Summary\|Metering\|Partial price | Shows whether the calculated price is partial for the resource |
| summary\|workload_indicator | Summary\|Workload Indicator | Workload Indicator |
| summary\|cpu_shares | Summary\|CPU Shares | CPU Shares |
| summary\|mem_shares | Summary\|Memory Shares | Memory Shares |
| summary\|number_datastore | Summary\|Number of Datastores | Number of Datastores |
| summary\|number_network | Summary\|Number of Networks | Number of Networks |
| guestfilesystem\|capacity | Guest File System\|Partition Capacity | Disk space capacity on guest file system partition. |
| guestfilesystem\|percentage | Guest File System\|Partition Utilization (%) | Guest file system partition space utilization in percentage |
| guestfilesystem\|usage | Guest File System\|Partition Utilization | Guest file system partition space utilization |
| guestfilesystem\|capacity_total | Guest File System\|Total Capacity | Disk space capacity on guest file system |
| guestfilesystem\|percentage_total | Guest File System\|Utilization (%) | Guest file system disk space utilization in percentage |
| guestfilesystem\|usage_total | Guest File System\|Utilization | Guest file system disk space utilization |
| guestfilesystem\|freespace | Guest File System\|Guest File System Free | Total free space on guest file system |
| guestfilesystem\|capacity_property | Guest File System\|Guest File System Capacity Property | Total capacity of guest file system as a property |
| guestfilesystem\|freespace_total | Guest File System\|Total Guest File System Free | Total free space on guest file system |
| guestfilesystem\| capacity_property_total | Guest File System\|Total Capacity Property | Total capacity of guest file system as a property |
| guest\|mem.free_latest | Guest\|Free Memory | Free Memory |
| guest\|mem.needed_latest | Guest\|Needed Memory | Needed Memory |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| guest\|mem.physUsable_latest | Guest\|Physically Usable Memory | Physically Usable Memory |
| guest\|page.inRate_latest | Guest\|Page In Rate per second | Page In Rate per second |
| guest\|page.size_latest | Guest\|Page Size | Page Size |
| guest\|swap.spaceRemaining_latest | Guest\|Remaining Swap Space | Remaining Swap Space |
| guest\|cpu_queue | Guest\|CPU Queue | The number of ready threads queuing in the CPU. Linux includes threads in running state. A number greater than 2 for prolong period indicates CPU core bottleneck. |
| guest\|disk_queue | Guest\|Disk Queue | The number of outstanding requests + IO currently in progress. |
| guest\|contextSwapRate_latest | Guest\|Context Swap Rate per second | Context Swap Rate per second |
| guest\|hugePage.size_latest | Guest\|Huge Page Size | Huge Page Size |
| guest\|hugePage.total_latest | Guest\|Total Huge Pages | Total Huge Pages |
| guest\|mem.activeFileCache_latest | Guest\|Active File Cache Memory | Active File Cache Memory |
| guest\|page.outRate_latest | Guest\|Page Out Rate per second | Page Out Rate per second |
| guest\|disk_queue_latest | Guest\|Disk Queue Latest | The number of outstanding requests + IO currently in progress. |
| virtualDisk\| numberReadAveraged_average | Virtual Disk\|Read IOPS | Number of read operations per second. This is averaged over the reporting period |
| virtualDisk\| numberWriteAveraged_average | Virtual Disk\|Write IOPS | Number of write operations per second. This is averaged over the reporting period |
| virtualDisk\|read_average | Virtual Disk\|Read Throughput | Amount of data read from storage in a second. This is averaged over the reporting period |
| virtualDisk\| totalReadLatency_average | Virtual Disk\|Read Latency | Average amount of time for a read operation by the storage adapter. |
| virtualDisk\| totalWriteLatency_average | Virtual Disk\|Write Latency | Average amount of time for a write operation by the storage adapter. |
| virtualDisk\|write_average | Virtual Disk\|Write Throughput | Amount of data written to storage in a second. This is averaged over the reporting period |
| virtualDisk\|usage | Virtual Disk\|Total Throughput | Amount of data read from/written to storage in a second. This is averaged over the reporting period |
| virtualDisk\|totalLatency | Virtual Disk\|Total Latency | Total Latency |
| virtualDisk\| commandsAveraged_average | Virtual Disk\|Total IOPS | Number of read/write operations per second. This is averaged over the reporting period |
| virtualDisk\|vDiskOIO | Virtual Disk\|Outstanding IO requests | OIO for datastore. |
| virtualDisk\|actualUsage | Virtual Disk\|Used Disk Space | Virtual Disk space usage |
| virtualDisk\|busResets_summation | Virtual Disk\|Bus Resets | The number of bus resets in the performance interval |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| virtualDisk\| commandsAborted_summation | Virtual Disk\|Commands Aborted | The number of disk commands stopped in the performance interval |
| virtualDisk\|readLoadMetric_latest | Virtual Disk\|Read Load | Storage DRS virtual disk metric read load |
| virtualDisk\|readOIO_latest | Virtual Disk\|Outstanding Read Requests | Average number of outstanding read requests to the virtual disk |
| virtualDisk\|writeLoadMetric_latest | Virtual Disk\|Write Load | Storage DRS virtual disk write load |
| virtualDisk\|writeOIO_latest | Virtual Disk\|Outstanding Write Requests | Average number of outstanding write requests to the virtual disk |
| virtualDisk\|smallSeeks_latest | Virtual Disk\|Number of Small Seeks | Small Seeks |
| virtualDisk\|mediumSeeks_latest | Virtual Disk\|Number of Medium Seeks | Medium Seeks |
| virtualDisk\|largeSeeks_latest | Virtual Disk\|Number of Large Seeks | Large Seeks |
| virtualDisk\|readLatencyUS_latest | Virtual Disk\|Read Latency (microseconds) | Read latency in microseconds |
| virtualDisk\|writeLatencyUS_latest | Virtual Disk\|Write Latency (microseconds) | Write Latency in microseconds |
| virtualDisk\|readIOSize_latest | Virtual Disk\|Average Read request size | Read IO size |
| virtualDisk\|writeIOSize_latest | Virtual Disk\|Average Write request size | Write IO size |
| diskspace\|pod_used | Disk Space\|Pod used | Space used by Pod files |
| diskspace\|provisionedSpace | Disk Space\|Provisioned Space for Pod | Provisioned space for Pod. In thin provisioned, it is the full space allocated (which may not be used yet). |
| diskspace\|notshared | Disk Space\|Not Shared | Space used by VM that is not shared with other VM |
| diskspace\|activeNotShared | Disk Space\|Active not shared | Unshared disk space used by VMs excluding snapshot |
| diskspace\|perDsUsed | Disk Space\|Pod used | Space used by all files of the Pod on the datastore (disks, snapshots, configs, logs, etc). |
| diskspace\|total_usage | Disk Space\|Utilization | Total disk space used on all datastores visible to this object |
| diskspace\|total_capacity | Disk Space\|Total Capacity | Total disk space on all datastores visible to this object |
| diskspace\|diskused | Disk Space\|Virtual Disk Used | Space used by virtual disks |
| diskspace\|snapshot | Disk Space\|Snapshot Space | Space used by snapshots |
| diskspace\|shared | Disk Space\|Shared Used | Shared space used |
| diskspace\|provisioned | Disk Space\|Provisioned Space | Provisioned space |
| diskspace\|snapshot\|used | Disk Space\|Snapshot\|Pod used | Disk space used by the Pod snapshot files. This is the space that can be potentially reclaimed if the snapshot is removed. |
| diskspace\|snapshot\|accessTime | Disk Space\|Snapshot\|Access Time | The date and time the snapshot was taken. |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| storage\|totalReadLatency_average | Storage\|Read Latency | Average amount of time for a read operation. |
| storage\|totalWriteLatency_average | Storage\|Write Latency | Average amount of time for a write operation. |
| storage\|read_average | Storage\|Read Rate | Read throughput rate |
| storage\|write_average | Storage\|Write Rate | Write throughput rate |
| storage\|usage_average | Storage\|Total Usage | Total throughput rate |
| storage\|numberReadAveraged_average | Storage\|Reads per second | Average number of read commands issued per second during the collection interval |
| storage\|numberWriteAveraged_average | Storage\|Writes per second | Average number of write commands issued per second during the collection interval |
| storage\|commandsAveraged_average | Storage\|Commands per second | Average number of commands issued per second during the collection interval |
| storage\|totalLatency_average | Storage\|Total Latency | Total latency |
| storage\|demandKBps | Storage\|Demand | N/A |
| storage\|contention | Storage\|Contention percentage | N/A |
| cost\|monthlyTotalCost | Cost\|MTD Total Cost | Month To Date Cost of Virtual Machine |
| cost\|monthlyProjectedCost | Cost\|Monthly Projected Total Cost | Virtual Machine cost projected for full month |
| cost\|compTotalCost | Cost\|MTD Compute Total Cost | Month to Date Total Compute Cost (Including CPU and Memory) of Virtual Machine |
| cost\|directCost | Cost\|Monthly Direct Cost | Monthly Direct Cost (comprising of OS Labor, VI Labor and any windows desktop instance license) of Virtual Machine |
| cost\|cpuCost | Cost\|MTD CPU Cost | Month to Date Virtual Machine CPU Cost. It is based on utilization. The more the VM uses, the higher its cost. |
| cost\|memoryCost | Cost\|MTD Memory Cost | Month to Date Memory Cost of Virtual Machine. It is based on utilization. The more the VM uses, the higher its cost. |
| cost\|storageCost | Cost\|MTD Disk Space Cost | Month to Date Disk Space Cost of Virtual Machine |
| cost\|reclaimableCost | Cost\|Potential Savings | Potential Savings |
| cost\|osLaborTotalCost | Cost\|Monthly OS Labor Cost | Operating System Labor Cost of Virtual Machine for full month |
| cost\|viLaborTotalCost | Cost\|Monthly VI Labor Cost | Monthly VI Labor Cost |
| cost\|effectiveTotalCost | Cost\|MTD Effective Total Cost | Month to Date Cost of Virtual Machine considering the allocation and demand model |
| cost\|effectiveProjectedTotalCost | Cost\|Monthly Effective Projected Total Cost | Virtual Machine cost projected for full month considering the allocation and demand model |
| cost\|allocation\|allocationBasedCpuMTDCost | Cost\|Allocation\|MTD CPU Cost | Month to Date Virtual Machine CPU Cost. It is based on utilization. The more the VM uses, the higher its cost. |
| cost\|allocation\|allocationBasedMemoryMTDCost | Cost\|Allocation\|MTD Memory Cost | Month to Date Memory Cost of Virtual Machine. It is based on utilization. The more the VM uses, the higher its cost. |
| cost\|allocation\|allocationBasedStorageMTDCost | Cost\|Allocation\|MTD Disk Space Cost | Month to Date Disk Space Cost of Virtual Machine |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| cost\|allocation\| allocationBasedTotalMTDCost | Cost\|Allocation\|MTD Total Cost | Month To Date Cost of Virtual Machine |
| cost\|allocation\| allocationBasedTotalCost | Cost\|Allocation\|Monthly Projected Total Cost | Virtual Machine cost projected for full month |
| datastore\|demand_oio | Datastore\|Outstanding IO requests | Amount of IO waiting in the queue to be executed. High IO, coupled with high latency, impacts performance. |
| datastore\| numberReadAveraged_average | Datastore\|Read IOPS | Number of read operations per second. This is averaged over the reporting period. |
| datastore\| numberWriteAveraged_average | Datastore\|Write IOPS | Number of write operations per second. This is averaged over the reporting period. |
| datastore\|read_average | Datastore\|Read Throughput | Amount of data read from storage in a second. This is averaged over the reporting period. |
| datastore\| totalReadLatency_average | Datastore\|Read Latency | Average amount of time for a read operation at the datastore level. It's an average of all the VMs in the datastore. |
| datastore\| totalWriteLatency_average | Datastore\|Write Latency | Average amount of time for a write operation by the storage adapter. |
| datastore\|write_average | Datastore\|Write Throughput | Amount of data written from storage in a second. This is averaged over the reporting period. |
| datastore\|totalLatency_average | Datastore\|Total Latency | Normalized Latency, taking into account the read/write ratio. |
| datastore\|usage_average | Datastore\|Total Throughput | Amount of data read from/written to storage in a second. This is averaged over the reporting period. |
| datastore\| commandsAveraged_average | Datastore\|Total IOPS | Number of read/write operations per second. This is averaged over the reporting period. |
| datastore\|used | Datastore\|Used Space | Used Space. |
| datastore\|demand | Datastore\|Demand | Max of datastore "Reads Per Sec", "Writes Per Sec", "Read Rate", "Write Rate", "OIO Per Sec" percentages. |
| datastore\|maxTotalLatency_latest | Datastore\|Highest Latency | Highest Latency. |
| datastore\|totalLatency_max | Datastore\|Total Latency Max | Total Latency Max (ms). |
| datastore\| maxObserved_NumberRead | Datastore\|Max Observed Reads per second | Max observed average number of read commands issued per second during the collection interval. |
| datastore\|maxObserved_Read | Datastore\|Max Observed Read Rate | Max observed rate of reading data from the datastore. |
| datastore\| maxObserved_NumberWrite | Datastore\|Max Observed Writes per second | Max observed average number of write commands issued per second during the collection interval. |
| datastore\|maxObserved_Write | Datastore\|Max Observed Write Rate | Max observed rate of writing data from the datastore. |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| datastore\|maxObserved_OIO | Datastore\|Max Observed Number of Outstanding IO Operations | N/A |

**Power Metrics for vSphere Pods**

Power metrics provide information about the vSphere pods power use.

| Metric Name | Description |
|---|---|
| Power\|Total Energy Consumed in the collection period (Wh) | Displays the total electricity consumed based on the time interval selected. The default collection cycle is set to 5 mins. You can continue using the default setting or edit it for each adapter instance. For example, if the time interval is set to its default value, the value represents the energy consumed per 5 mins. |
| Power\|Current Power Consumption Rate (Watt) | The power consumption rate per second, averaged over the reporting period.<br>Key: power\|power_average |
| Power\|(DEP) Energy (Joule) | Total energy consumed in joules.<br>Key: power\|energy_summation |

### OS and Application Monitoring Metrics

Metrics are collected for operating systems, application services, remote checks, Linux processes, and Windows services.

### Operating System Metrics

Metrics are collected for Linux and Windows operating systems.

**Linux Platforms**

The following metrics are collected for Linux operating systems:

**Table 418: Metrics for Linux**

| Metric | Metric Category | KPI |
|---|---|---|
| <Instance name>\| Usage Idle | CPU | False |
| <Instance name>\| Usage IO-Wait | CPU | False |
| <Instance name>\|Time Active | CPU | True |
| <Instance name>\|Time Guest | CPU | False |
| <Instance name>\|Time Guest Nice | CPU | False |
| <Instance name>\|Time Idle | CPU | False |
| <Instance name>\|Time IO-Wait | CPU | False |

*Table continued on next page*

*Continued from previous page*

| Metric | Metric Category | KPI |
|---|---|---|
| <Instance name>\|Time IRQ | CPU | True |
| <Instance name>\|Time Nice | CPU | False |
| <Instance name>\|Time Soft IRQ | CPU | True |
| <Instance name>\|Time Steal | CPU | False |
| <Instance name>\|Time System | CPU | False |
| <Instance name>\|Time User | CPU | True |
| <Instance name>\|Usage Active (%) | CPU | True |
| <Instance name>\|Usage Guest (%) | CPU | False |
| <Instance name>\|Usage Guest Nice (%) | CPU | False |
| <Instance name>\|Usage IRQ (%) | CPU | True |
| <Instance name>\|Usage Nice (%) | CPU | False |
| <Instance name>\|Usage Soft IRQ (%) | CPU | True |
| <Instance name>\|Usage Steal (%) | CPU | False |
| <Instance name>\|Usage System (%) | CPU | True |
| <Instance name>\|Usage User (%) | CPU | True |
| CPU Load1 (%) | CPU Load | False |
| CPU Load15 (%) | CPU Load | False |
| CPU Load5 (%) | CPU Load | False |
| <Instance name>\|IO Time | Disk IO | False |
| <Instance name>\|Read Time | Disk IO | False |
| <Instance name>\|Reads | Disk IO | False |
| <Instance name>\|Write Time | Disk IO | False |
| <Instance name>\|Writes | Disk IO | False |
| <Instance name>\|Disk Free | Disk | False |
| <Instance name>\|Disk Total | Disk | False |
| <Instance name>\|Disk Used (%) | Disk | False |
| Cached | Memory | False |
| Free | Memory | False |
| Inactive | Memory | False |
| Total | Memory | True |
| Used | Memory | True |
| Used Percent | Memory | True |
| Blocked | Processes | True |
| Dead | Processes | False |
| Running | Processes | False |
| Sleeping | Processes | False |
| Stopped | Processes | False |
| Zombies | Processes | False |
| Free | Swap | False |
| In | Swap | False |
| Out | Swap | False |

*Table continued on next page*

*Continued from previous page*

| Metric | Metric Category | KPI |
|---|---|---|
| Total | Swap | True |
| Used | Swap | True |
| Used Percent | Swap | True |
| Telegraf Availability | None | False |

**Windows Platforms**

The following metrics are collected for Windows operating systems:

**Table 419: Metrics for Windows**

| Metric | Metric Category | KPI |
|---|---|---|
| Idle Time | CPU | False |
| Interrupt Time | CPU | False |
| Interrupts persec | CPU | True |
| Privileged Time | CPU | False |
| Processor Time | CPU | False |
| User Time | CPU | False |
| DPC Time (%) | CPU | False |
| Usage Guest (%) | CPU | False |
| Usage System (%) | CPU | False |
| Usage User (%) | CPU | False |
| Avg. Disk Bytes Read | Disk | False |
| Avg. Disk sec Read | Disk | False |
| Avg. Disk sec Write | Disk | False |
| Avg. Disk Write Queue Length | Disk | False |
| Avg. Disk Read Queue Length | Disk | False |
| Disk Read Time | Disk | False |
| Disk Write Time | Disk | False |
| Free Megabytes | Disk | False |
| Free Space | Disk | False |
| Idle Time | Disk | False |
| Split IO persec | Disk | False |
| Available Bytes | Memory | True |
| Cache Bytes | Memory | False |
| Cache Faults persec | Memory | False |
| Committed Bytes | Memory | True |
| Demand Zero Faults persec | Memory | False |
| Page Faults persec | Memory | True |
| Pages persec | Memory | False |
| Pool Nonpaged Bytes | Memory | True |
| Pool Paged Bytes | Memory | False |

*Table continued on next page*

*Continued from previous page*

| Metric | Metric Category | KPI |
|---|---|---|
| Transition Faults persec | Memory | False |
| Total (bytes) | Memory | False |
| Used (bytes) | Memory | False |
| Used Percent(%) | Memory | False |
| Bytes Received persec | Network | False |
| Bytes Sent persec | Network | False |
| Packets Outbound Discarded | Network | False |
| Packets Outbound Errors | Network | False |
| Packets Received Discarded | Network | False |
| Packets Received Errors | Network | False |
| Packets Received persec | Network | False |
| Packets Sent persec | Network | False |
| Elapsed Time | Process | False |
| Handle Count | Process | False |
| IO Read Bytes persec | Process | False |
| IO Read Operations persec | Process | False |
| IO Write Bytes persec | Process | False |
| IO Write Operations persec | Process | False |
| Privileged Time | Process | False |
| Processor Time | Process | False |
| Thread Count | Process | False |
| User Time | Process | False |
| Context Switches persec | System | False |
| Processes | System | False |
| Processor Queue Length | System | False |
| System Calls persec | System | False |
| System Up Time | System | False |
| Threads | System | False |
| Used Percent (%) | Swap | False |
| Total (bytes) | Swap | False |
| Telegraf Availability | None | False |

### Application Service Metrics

Metrics are collected for 23 application services.

### Active Directory Metrics

Metrics are collected for the Active Directory application service.

**Table 420: Active Directory Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Database Cache % Hit (%) | Active Directory Database | True |
| Database Cache Page Faults/sec | Active Directory Database | True |
| Database Cache Size | Active Directory Database | False |
| Data Lookups | Active Directory DFS Replication | False |
| Database Commits | Active Directory DFS Replication | True |
| Avg Response Time | Active Directory DFSN | True |
| Requests Failed | Active Directory DFSN | False |
| Requests Processed | Active Directory DFSN | False |
| Dynamic Update Received | Active Directory DNS | False |
| Dynamic Update Rejected | Active Directory DNS | False |
| Recursive Queries | Active Directory DNS | False |
| Recursive Queries Failure | Active Directory DNS | False |
| Secure Update Failure | Active Directory DNS | False |
| Total Query Received | Active Directory DNS | True |
| Total Response Sent | Active Directory DNS | True |
| Digest Authentications | Active Directory Security System-Wide Statistics | True |
| Kerberos Authentications | Active Directory Security System-Wide Statistics | True |
| NTLM Authentications | Active Directory Security System-Wide Statistics | True |
| Directory Services:<InstanceName>\| Base Searches persec | Active Directory Services | False |
| Directory Services:<InstanceName>\| Database adds persec | Active Directory Services | False |
| Directory Services:<InstanceName>\| Database deletes persec | Active Directory Services | False |
| Directory Services<InstanceName>\| Database modifys/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| Database recycles/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Inbound Bytes Total/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Inbound Objects/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Outbound Bytes Total/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Outbound Objects/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Pending Replication Operations | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Pending Replication Synchronizations | Active Directory Services | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Directory Services<InstanceName>\| DRA Sync Requests Made | Active Directory Services | False |
| Directory Services<InstanceName>\| DRA Sync Requests Successful | Active Directory Services | False |
| Directory Services<InstanceName>\| DS Client Binds/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| DS Directory Reads/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| DS Directory Searches/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| DS Server Binds/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| DS Threads in Use | Active Directory Services | True |
| Directory Services:<InstanceName>\| LDAP Active Threads | Active Directory Services | False |
| Directory Services:<InstanceName>\| LDAP Client Sessions | Active Directory Services | True |
| Directory Services<InstanceName>\| LDAP Closed Connections/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| LDAP New Connections/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| LDAP Searches/sec | Active Directory Services | True |
| Directory Services<InstanceName>\| LDAP Successful Binds/sec | Active Directory Services | False |
| Directory Services<InstanceName>\| LDAP UDP operations/sec | Active Directory Services | False |
| Directory Services:<InstanceName>\| LDAP Writes/sec | Active Directory Services | False |
| Application Availability | Active Directory | False |

### ActiveMQ Metrics

Metrics are collected for the ActiveMQ application service.

**Table 421: ActiveMQ Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Active MQ | False |
| Buffer Pool<InstanceName>\| Memory Used | Active MQ | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Total Capacity | Active MQ | False |
| Class Loading\|Loaded Class Count | Active MQ | False |
| Class Loading\|Unloaded Class Count | Active MQ | False |
| Class Loading\|Total Loaded Class Count | Active MQ | False |
| File Descriptor Usage\|Max File Descriptor Count | Active MQ | False |
| File Descriptor Usage\|Open File Descriptor Count | Active MQ | False |
| Garbage Collection<InstanceName>\| Total Collection Count | Active MQ | False |
| Garbage Collection<InstanceName>\| Total Collection Time | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Committed Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Initial Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Maximum Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Used Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Committed Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Initial Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Maximum Memory | Active MQ | False |
| JVM Memory Pool<InstanceName>\| Usage\|Used Memory | Active MQ | False |
| Application Availability | Active MQ | False |
| Threading\|Thread Count | Active MQ | False |
| Uptime | Active MQ | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Process CpuLoad | Active MQ | False |
| UTILIZATION\|Memory Limit | ActiveMQ Broker | True |
| UTILIZATION\|Memory Percent Usage (%) | ActiveMQ Broker | True |
| UTILIZATION\|Store Limit | ActiveMQ Broker | False |
| UTILIZATION\|Store Percent Usage (%) | ActiveMQ Broker | False |
| UTILIZATION\|Temp Limit | ActiveMQ Broker | False |
| UTILIZATION\|Temp Percent Usage (%) | ActiveMQ Broker | False |
| UTILIZATION\|Total Consumer Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Dequeue Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Enqueue Count | ActiveMQ Broker | True |
| UTILIZATION\|Total Message Count | ActiveMQ Broker | True |
| JVM Memory\|Heap Memory Usage\|Initial Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Maximum Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Heap Memory Usage\|Used Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | ActiveMQ JVM Memory Usage | False |
| JVM Memory\|Object Pending FinalizationCount | ActiveMQ JVM Memory Usage | False |
| UTILIZATION\|Process CpuLoad | ActiveMQ OS | False |
| UTILIZATION\|System Cpu Load | ActiveMQ OS | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Consumer Count | ActiveMQ Topic | True |
| UTILIZATION\|Dequeue Count | ActiveMQ Topic | True |
| UTILIZATION\|Enqueue Count | ActiveMQ Topic | True |
| UTILIZATION\|Queue Size | ActiveMQ Topic | True |
| UTILIZATION\|Producer Count | ActiveMQ Topic | False |

### Apache HTTPD Metrics

Metrics are collected for the Apache HTTPD application service.

**NOTE**
Metrics are collected for the Events MPM. Metrics are not collected for the other MPMs.

**Table 422: Apache HTTPD Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Busy Workers | Apache HTTPD | True |
| UTILIZATION\|Bytes Per Req | Apache HTTPD | False |
| UTILIZATION\|Bytes Per Sec | Apache HTTPD | False |
| UTILIZATION\|CPU Load | Apache HTTPD | True |
| UTILIZATION\|CPU User | Apache HTTPD | False |
| UTILIZATION\|Idle Workers | Apache HTTPD | True |
| UTILIZATION\|Request Per Sec | Apache HTTPD | True |
| UTILIZATION\|SCBoard Closing | Apache HTTPD | False |
| UTILIZATION\|SCBoard DNS Lookup | Apache HTTPD | False |
| UTILIZATION\|SCBoard Finishing | Apache HTTPD | False |
| UTILIZATION\|SCBoard Idle Cleanup | Apache HTTPD | False |
| UTILIZATION\|SCBoard Keep Alive | Apache HTTPD | False |
| UTILIZATION\|SCBoard Logging | Apache HTTPD | False |
| UTILIZATION\|SCBoard Open | Apache HTTPD | False |
| UTILIZATION\|SCBoard Reading | Apache HTTPD | False |
| UTILIZATION\|SCBoard Sending | Apache HTTPD | False |
| UTILIZATION\|SCBoard Starting | Apache HTTPD | False |
| UTILIZATION\|SCBoard Waiting | Apache HTTPD | False |
| UTILIZATION\|Total Accesses | Apache HTTPD | False |
| UTILIZATION\|Total Bytes | Apache HTTPD | True |
| UTILIZATION\|Total Connections | Apache HTTPD | False |
| UTILIZATION\|Uptime | Apache HTTPD | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Asynchronous Closing Connections | Apache HTTPD | False |
| UTILIZATION\|Asynchronous Keep Alive Connections | Apache HTTPD | False |
| UTILIZATION\|Asynchronous Writing Connections | Apache HTTPD | False |
| UTILIZATION\|ServerUptimeSeconds | Apache HTTPD | False |
| UTILIZATION\|Load1 | Apache HTTPD | False |
| UTILIZATION\|Load5 | Apache HTTPD | False |
| UTILIZATION\| ParentServerConfigGeneration | Apache HTTPD | False |
| UTILIZATION\| ParentServerMPMGeneration | Apache HTTPD | False |
| Application Availability | Apache HTTPD | False |

**Apache HTTPD**

Metrics are collected for the Apache HTTPD application service.

**Table 423: Apache Tomcat**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Tomcat Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Tomcat Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Tomcat Server | False |
| Class Loading\|Loaded Class Count | Tomcat Server | False |
| Class Loading\|Total Loaded Class Count | Tomcat Server | False |
| Class Loading\|Unloaded Class Count | Tomcat Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Tomcat Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Tomcat Server | False |
| Garbage Collection:<InstanceName>\|Total Collection Count | Tomcat Server | False |
| Garbage Collection:<InstanceName>\|Total Collection Time | Tomcat Server | True |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Tomcat Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Heap Memory Usage\| Initial Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Tomcat Server | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Tomcat Server | False |
| JVM Memory\|Number of Object Pending Finalization Count | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Used Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Committed Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Initial Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Maximum Memory | Tomcat Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Used Memory | Tomcat Server | False |
| Process CPU Usage (%) | Tomcat Server | True |
| System CPU Usage (%) | Tomcat Server | True |
| System Load Average (%) | Tomcat Server | True |
| Threading\|Thread Count | Tomcat Server | False |
| Uptime | Tomcat Server | True |
| Application Availability | Tomcat Server | False |
| JSP Count | Tomcat Server Web Module | False |
| JSP Reload Count | Tomcat Server Web Module | False |
| JSP Unload Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Count | Tomcat Server Web Module | False |
| Servlet:<InstanceName>\|Total Request Error Count | Tomcat Server Web Module | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Servlet:<InstanceName>\|Total Request Processing Time | Tomcat Server Web Module | False |
| Cache : Hit Count | Tomcat Server Web Module | False |
| Cache : Lookup Count | Tomcat Server Web Module | False |
| Current Thread Count | Tomcat Server Global Request Processor | True |
| Current Threads Busy | Tomcat Server Global Request Processor | True |
| errorRate | Tomcat Server Global Request Processor | False |
| Total Request Bytes Received | Tomcat Server Global Request Processor | False |
| Total Request Bytes Sent | Tomcat Server Global Request Processor | False |
| Total Request Count | Tomcat Server Global Request Processor | True |
| Total Request Error Count | Tomcat Server Global Request Processor | True |
| Total Request Processing Time | Tomcat Server Global Request Processor | False |

### Microsoft IIS Metrics

Metrics are collected for the Microsoft IIS application service.

**Table 424: IIS Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Service Request Queues<InstanceName>AppPool\| CurrentQueueSize | IIS HTTP Service Request Queues | True |
| HTTP Service Request Queues<InstanceName>AppPool\| RejectedRequests | IIS HTTP Service Request Queues | False |
| Web Services<InstanceName> Web Site\|Bytes Received | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Bytes Sent/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Bytes Total/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Connection Attempts/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Current Connections | IIS Web Services | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Web Services<InstanceName> Web Site\|Get Requests/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Locked Errors/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Not Found Errors/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Post Requests/sec | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Service Uptime | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Total Bytes Sent | IIS Web Services | False |
| Web Services<InstanceName> Web Site\|Total Get Requests | IIS Web Services | True |
| Web Services<InstanceName> Web Site\|Total Post Requests | IIS Web Services | True |
| Web Services<InstanceName> Web Site\|Total Put Requests | IIS Web Services | False |
| Current File Cache Memory Usage (bytes) | IIS Web Services Cache | False |
| File Cache Hits Percent (%) | IIS Web Services Cache | False |
| Kernel URI Cache Hits Percent (%) | IIS Web Services Cache | False |
| Kernel URI Cache Misses | IIS Web Services Cache | False |
| Total Flushed URIs | IIS Web Services Cache | False |
| URI Cache Hits | IIS Web Services Cache | False |
| URI Cache Hits Percent (%) | IIS Web Services Cache | False |
| URI Cache Misses | IIS Web Services Cache | False |
| ASP.NET<InstanceName>\|Application Restarts | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Request Wait Time | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Requests Current | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Requests Queued | IIS ASP.NET | True |
| ASP.NET<InstanceName>\|Requests Rejected | IIS ASP.NET | True |
| MS.NET<InstanceName>\|Allocated Bytes/sec | MS.NET | True |
| MS.NET<InstanceName>\|Current Queue Length | MS.NET | False |
| MS.NET<InstanceName>\|Finalization Survivors | MS.NET | False |
| MS.NET<InstanceName>\|Gen 0 Collections | MS.NET | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| MS.NET<InstanceName>\|Gen 0 heap size | MS.NET | False |
| MS.NET<InstanceName>\|Gen 1 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 1 heap size | MS.NET | False |
| MS.NET<InstanceName>\|Gen 2 Collections | MS.NET | False |
| MS.NET<InstanceName>\|Gen 2 heap size | MS.NET | False |
| MS.NET<InstanceName>\|IL Bytes Jitted / sec | MS.NET | False |
| MS.NET<InstanceName>\|Induced GC | MS.NET | False |
| MS.NET<InstanceName>\|Large Object Heap size | MS.NET | False |
| MS.NET<InstanceName>\|No of current logical Threads | MS.NET | True |
| MS.NET<InstanceName>\|No of current physical Threads | MS.NET | True |
| MS.NET<InstanceName>\|No of current recognized threads | MS.NET | False |
| MS.NET<InstanceName>\|No of Exceps Thrown / sec | MS.NET | True |
| MS.NET<InstanceName>\|No of total recognized threads | MS.NET | False |
| MS.NET<InstanceName>\|Percent Time in Jit | MS.NET | False |
| MS.NET<InstanceName>\|Pinned Objects | MS.NET | False |
| MS.NET<InstanceName>\|Stack Walk Depth | MS.NET | False |
| MS.NET<InstanceName>\|Time in RT checks | MS.NET | False |
| MS.NET<InstanceName>\|Time Loading | MS.NET | True |
| MS.NET<InstanceName>\|Total No of Contentions | MS.NET | False |
| MS.NET<InstanceName>\|Total Runtime Checks | MS.NET | True |
| Application Availability | Microsoft IIS | False |

### Java Application Metrics

Metrics are collected for the Java application service.

**Table 425: Java Application Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Java Application | False |
| Buffer Pool<InstanceName>\|Memory Used | Java Application | False |
| Buffer Pool<InstanceName>\|Total Capacity | Java Application | False |
| Class Loading\|Loaded Class Count | Java Application | True |
| Class Loading\|Total Loaded Class Count | Java Application | False |
| Class Loading\|Unloaded Class Count | Java Application | False |
| Garbage Collection<InstanceName>\| Total Collection Count | Java Application | False |
| Garbage Collection<InstanceName>\| Total Collection Time | Java Application | False |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Java Application | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Committed Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Initial Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Maximum Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Peak Usage\| Used Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\| Committed Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Initial Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\| Maximum Memory | Java Application | False |
| JVM Memory\|JVM Memory Pool<InstanceName>\|Usage\|Used Memory | Java Application | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Java Application | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Java Application | False |
| JVM Memory\|Object Pending Finalization Count | Java Application | False |
| Uptime | Java Application | True |
| Threading\|Thread Count | Java Application | True |
| Process CPU Usage % | Java Application | False |
| System CPU Usage % | Java Application | False |
| System Load Average % | Java Application | False |

### JBoss Server Metrics

Metrics are collected for the JBoss Server application service.

**Table 426: JBoss Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Jboss Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Jboss Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Jboss Server | False |
| Class Loading\|Loaded Class Count | Jboss Server | False |
| Class Loading\|Total Loaded Class Count | Jboss Server | False |
| Class Loading\|Unloaded Class Count | Jboss Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Jboss Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Jboss Server | False |
| Http Listener<InstanceName>\|Bytes Received | Jboss Server | False |
| Http Listener<InstanceName>\|Bytes Sent | Jboss Server | False |
| Http Listener<InstanceName>\|Error Count | Jboss Server | False |
| Http Listener<InstanceName>\| Request Count | Jboss Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Https Listener<InstanceName>\|Bytes Received | Jboss Server | False |
| Https Listener<InstanceName>\|Bytes Sent | Jboss Server | False |
| Https Listener<InstanceName>\|Error Count | Jboss Server | False |
| Https Listener<InstanceName>\| Request Count | Jboss Server | False |
| Process CPU Usage (%) | Jboss Server | False |
| System CPU Usage (%) | Jboss Server | False |
| System Load Average (%) | Jboss Server | False |
| Threading\|Daemon Thread Count | Jboss Server | False |
| Threading\|Peak Thread Count | Jboss Server | False |
| Threading\|Thread Count | Jboss Server | False |
| Threading\|Total Started Thread Count | Jboss Server | False |
| Uptime | Jboss Server | False |
| UTILIZATION\|Heap Memory Usage | Jboss Server | False |
| Application Availability | Jboss Server | False |
| Garbage Collection<InstanceName>\| Total Collection Count | Jboss JVM Garbage Collector | False |
| Garbage Collection<InstanceName>\| Total Collection Time | Jboss JVM Garbage Collector | False |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Jboss JVM Memory | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Jboss JVM Memory | True |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Jboss JVM Memory | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Jboss JVM Memory | False |
| JVM Memory\|Object Pending Finalization Count | Jboss JVM Memory | True |
| UTILIZATION\|Active Count | Jboss Datasource Pool | False |
| UTILIZATION\|Available Count | Jboss Datasource Pool | False |
| JVM Memory Pool<InstanceName>\| Collection Usage\|Committed Memory | Jboss JVM Memory Pool | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory Pool<InstanceName>\|Collection Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Collection Usage\|Used Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Collection Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Peak Usage\|Used Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Usage\|Committed Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Usage\|Initial Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Usage\|Maximum Memory | Jboss JVM Memory Pool | False |
| JVM Memory Pool<InstanceName>\|Usage\|Used Memory | Jboss JVM Memory Pool | False |

### HyperV Metrics

Metrics are collected for the HyperV application service.

**Table 427: HyperV Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| VM:Hyper-V Virtual Machine Health Summary\|Health Critical | HyperV | False |
| VM<instanceName>\|Physical Memory | HyperV | False |
| VM<instanceName>Hv VP 0\|Total Run Time | HyperV | False |
| VM<instanceName>\|Bytes Received | HyperV | False |
| VM<instanceName>\|Bytes Sent | HyperV | False |
| VM<instanceName>\|Error Count | HyperV | False |
| VM<instanceName>\|Latency | HyperV | False |
| VM<instanceName>\|Queue Length | HyperV | False |
| VM<instanceName>\|Throughput | HyperV | False |
| CPU<instanceName>\|Idle Time | HyperV | True |
| CPU<instanceName>\|Processor Time | HyperV | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| CPU<instanceName>\|User Time | HyperV | True |
| Disk<instanceName>\|Avg Disk Queue Length | HyperV | False |
| Disk<instanceName>\|Idle Time | HyperV | False |
| Disk<instanceName>\|Read Time | HyperV | True |
| Disk<instanceName>\|Write Time | HyperV | True |
| Process<instanceName>\|Private Bytes | HyperV | False |
| Process<instanceName>\|Processor Time | HyperV | False |
| Process<instanceName>\|Thread Count | HyperV | False |
| Process<instanceName>\|User Time | HyperV | False |
| System\|Processes | HyperV | False |
| System\|Processor Queue Length | HyperV | False |
| System\|System UpTime | HyperV | False |
| Memory\|Available Bytes | HyperV | False |
| Memory\|Cache Bytes | HyperV | False |
| Memory\|Cache Faults | HyperV | False |
| Memory\|Pages | HyperV | False |
| Network<instanceName>\|Packets Outbound Error | HyperV | False |
| Network<instanceName>\|Packets Received Error | HyperV | False |
| Application Availability | HyperV | False |

### Oracle DB Metrics

Metrics are collected for the Oracle DB application service.

Oracle DB cannot be activated on Linux platforms.

**Table 428: Oracle DB Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Utilization\|Active Sessions | OracleDB | True |
| Utilization\|Buffer CacheHit Ratio | OracleDB | False |
| Utilization\|Cursor CacheHit Ratio | OracleDB | False |
| Utilization\|Database Wait Time | OracleDB | False |
| Utilization\|Disk Sort persec | OracleDB | False |
| Utilization\|Enqueue Timeouts Persec | OracleDB | False |
| Utilization\|Global Cache Blocks Corrupted | OracleDB | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Utilization\|Global Cache Blocks Lost | OracleDB | False |
| Utilization\|Library CacheHit Ratio | OracleDB | False |
| Utilization\|Logon persec | OracleDB | True |
| Utilization\|Memory Sorts Ratio | OracleDB | True |
| Utilization\|Rows persort | OracleDB | False |
| Utilization\|Service Response Time | OracleDB | False |
| Utilization\|Session Count | OracleDB | True |
| Utilization\|Session Limit | OracleDB | False |
| Utilization\|Shared Pool Free | OracleDB | False |
| Utilization\|Temp Space Used | OracleDB | False |
| Utilization\|Total Sorts persec | OracleDB | False |
| Utilization\|Physical Read Bytes Persc | OracleDB | False |
| Utilization\|Physical Read IO Requests Persc | OracleDB | False |
| Utilization\|Physical Read Total Bytes Persec | OracleDB | False |
| Utilization\|Physical Reads Persec | OracleDB | True |
| Utilization\|Physical Reads Per Txn | OracleDB | False |
| Utilization\|Physical Write Bytes Persc | OracleDB | False |
| Utilization\|Physical Write IO Requests Persc | OracleDB | False |
| Utilization\|Physical Write Total Bytes Persc | OracleDB | False |
| Utilization\|Physical Writes Persc | OracleDB | True |
| Utilization\|Physical Writes Per Txn | OracleDB | False |
| Utilization\|User Commits Percentage | OracleDB | False |
| Utilization\|User Commits Persc | OracleDB | False |
| Utilization\|User Rollbacks Percentage | OracleDB | False |
| Utilization\|User Rollbacks persec | OracleDB | True |
| Utilization\|User Transaction Persec | OracleDB | False |
| Utilization\|Database Time Persc | OracleDB | False |
| Application Availability | Oracle DB | False |

### Cassandra Metrics

Metrics are collected for the Cassandra application service.

**Table 429: Cassandra Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Cache<InstanceName>\|Capacity | Cassandra | False |
| Cache<InstanceName>\|Entries | Cassandra | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Cache<InstanceName>\|HitRate | Cassandra | True |
| Cache<InstanceName>\|Requests | Cassandra | True |
| Cache<InstanceName>\|Size | Cassandra | False |
| ClientRequest<InstanceName>\| Failures | Cassandra | False |
| ClientRequest<InstanceName>\| Latency | Cassandra | False |
| ClientRequest<InstanceName>\| Timeouts | Cassandra | False |
| ClientRequest<InstanceName>\|Total Latency | Cassandra | False |
| ClientRequest<InstanceName>\| Unavailables | Cassandra | False |
| CommitLog\|Pending Tasks | Cassandra | False |
| CommitLog\|Total Commit Log Size | Cassandra | False |
| Compaction\|Bytes Compacted | Cassandra | False |
| Compaction\|Completed Tasks | Cassandra | False |
| Compaction\|Pending Tasks | Cassandra | False |
| Compaction\|Total Compactions Completed | Cassandra | False |
| Connected Native Clients | Cassandra | False |
| HeapMemoryUsage\|committed | Cassandra | False |
| HeapMemoryUsage\|init | Cassandra | False |
| HeapMemoryUsage\|max | Cassandra | False |
| HeapMemoryUsage\|used | Cassandra | False |
| NonHeapMemoryUsage\|committed | Cassandra | False |
| NonHeapMemoryUsage\|init | Cassandra | False |
| NonHeapMemoryUsage\|max | Cassandra | False |
| NonHeapMemoryUsage\|used | Cassandra | False |
| ObjectPendingFinalizationCount | Cassandra | False |
| Storage\|Exceptions Count | Cassandra | False |
| Storage\|Load Count | Cassandra | False |
| Table<InstanceName>\|Coordinator Read Latency | Cassandra | False |
| Table<InstanceName>\|Live Diskspace Used | Cassandra | False |
| Table<InstanceName>\|Read Latency | Cassandra | False |
| Table<InstanceName>\|Total Diskspace Used | Cassandra | False |
| Table<InstanceName>\|Total Read Latency | Cassandra | False |
| Table<InstanceName>\|Total Write Latency | Cassandra | False |
| Table<InstanceName>\|Write Latency | Cassandra | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| ThreadPools<InstanceName>\|Active Tasks | Cassandra | False |
| ThreadPools<InstanceName>\| Currently Blocked Tasks | Cassandra | False |
| ThreadPools<InstanceName>\|Pending Tasks | Cassandra | False |
| Application Availability | Cassandra | False |

### MongoDB Metrics

Metrics are collected for the MongoDB application service.

**Table 430: MongoDB Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Active Reads | MongoDB | True |
| UTILIZATION\|Active Writes | MongoDB | True |
| UTILIZATION\|Connections Available | MongoDB | False |
| UTILIZATION\|Connections Total Created | MongoDB | False |
| UTILIZATION\|Current Connections | MongoDB | True |
| UTILIZATION\|Cursor Timed Out | MongoDB | True |
| UTILIZATION\|Deletes Per Sec | MongoDB | False |
| UTILIZATION\|Document Inserted | MongoDB | False |
| UTILIZATION\|Document Deleted | MongoDB | False |
| UTILIZATION\|Flushes Per Sec | MongoDB | False |
| UTILIZATION\|Inserts Per Sec | MongoDB | False |
| UTILIZATION\|Net Input Bytes | MongoDB | False |
| UTILIZATION\|Open Connections | MongoDB | True |
| UTILIZATION\|Page Faults Per Second | MongoDB | False |
| UTILIZATION\|Net Output Bytes | MongoDB | False |
| UTILIZATION\|Queries Per Sec | MongoDB | False |
| UTILIZATION\|Queued Reads | MongoDB | True |
| UTILIZATION\|Queued Writes | MongoDB | True |
| UTILIZATION\|Total Available | MongoDB | False |
| UTILIZATION\|Total Deletes Per Sec | MongoDB | False |
| UTILIZATION\|Total Passes Per Sec | MongoDB | False |
| UTILIZATION\|Total Refreshing | MongoDB | False |
| UTILIZATION\|Updates Per Sec | MongoDB | False |
| UTILIZATION\|Volume Size MB | MongoDB | False |
| Application Availability | MongoDB | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Collection Stats | MongoDB DataBases | False |
| UTILIZATION\|Data Index Stats | MongoDB DataBases | True |
| UTILIZATION\|Data Indexes | MongoDB DataBases | False |
| UTILIZATION\|Data Size Stats | MongoDB DataBases | True |
| UTILIZATION\|Average Object Size stats | MongoDB DataBases | False |
| UTILIZATION\|Num Extents Stats | MongoDB DataBases | False |

### MS Exchange Metrics

Metrics are collected for the MS Exchange application service.

**Table 431: MS Exchange Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Active Manager Server\|Active Manager Role | MS Exchange | False |
| Active Manager Server\|Database State Info Writes per second | MS Exchange | False |
| Active Manager Server\| GetServerForDatabase Server-Side Calls | MS Exchange | False |
| Active Manager Server\|Server-Side Calls per second | MS Exchange | True |
| Active Manager Server\|Total Number of Databases | MS Exchange | True |
| ActiveSync\|Average Request Time | MS Exchange | True |
| ActiveSync\|Current Requests | MS Exchange | False |
| ActiveSync\|Mailbox Search Total | MS Exchange | False |
| ActiveSync\|Ping Commands Pending | MS Exchange | False |
| ActiveSync\|Requests per second | MS Exchange | True |
| ActiveSync\|Sync Commands per second | MS Exchange | True |
| ASP.NET\|Application Restarts | MS Exchange | False |
| ASP.NET\|Request Wait Time | MS Exchange | True |
| ASP.NET\|Worker Process Restarts | MS Exchange | False |
| Autodiscover Service\|Requests per second | MS Exchange | True |
| Availability Service\|Average Time to Process a Free Busy Request | MS Exchange | True |
| Outlook Web Access\|Average Search Time | MS Exchange | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Outlook Web Access\|Requests per second | MS Exchange | False |
| Outlook Web Access\|Current Unique Users | MS Exchange | False |
| Application Availability | MS Exchange | False |
| Performance\|Database Cache Hit (%) | MS Exchange Database | False |
| Performance\|Database Page Fault Stalls per second | MS Exchange Database | True |
| Performance\|I/O Database Reads Average Latency | MS Exchange Database | True |
| Performance\|I/O Database Writes Average Latency | MS Exchange Database | True |
| Performance\|I/O Log Reads Average Latency | MS Exchange Database | False |
| Performance\|I/O Log Writes Average Latency | MS Exchange Database | False |
| Performance\|Log Record Stalls per second | MS Exchange Database | False |
| Performance\|Log Threads Waiting | MS Exchange Database | False |
| Performance\|I/O Database Reads Average Latency | MS Exchange Database Instance | False |
| Performance\|I/O Database Writes Average Latency | MS Exchange Database Instance | False |
| Performance\|Log Record Stalls per second | MS Exchange Database Instance | False |
| Performance\|Log Threads Waiting | MS Exchange Database Instance | False |
| Performance\|LDAP Read Time | MS Exchange Domain Controller | False |
| Performance\|LDAP Search Time | MS Exchange Domain Controller | False |
| Performance\|LDAP Searches Timed Out per minute | MS Exchange Domain Controller | False |
| Performance\|Long Running LDAP Operations per minute | MS Exchange Domain Controller | False |
| Performance\|Connection Attempts per second | MS Exchange Web Server | True |
| Performance\|Current Connections | MS Exchange Web Server | False |
| Performance\|Other Request Methods per second | MS Exchange Web Server | False |
| Process\|Handle Count | MS Exchange Windows Service | False |
| Process\|Memory Allocated | MS Exchange Windows Service | False |
| Process\|Processor Time (%) | MS Exchange Windows Service | True |
| Process\|Thread Count | MS Exchange Windows Service | False |
| Process\|Virtual Memory Used | MS Exchange Windows Service | False |
| Process\|Working Set | MS Exchange Windows Service | False |

### Microsoft SQL Server Metrics

Metrics are collected for the Microsoft SQL Server application service.

**Table 432: MS SQL Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| CPU<InstanceName>\|CPU Usage (%) | Microsoft SQL Server | False |
| Database IO\|Rows Reads Bytes/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Reads/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Writes Bytes/Sec | Microsoft SQL Server | False |
| Database IO\|Rows Writes/Sec | Microsoft SQL Server | False |
| Performance\|Access Methods\|Full Scans per second | Microsoft SQL Server | False |
| Performance\|Access Methods\|Index Searches | Microsoft SQL Server | False |
| Performance\|Access Methods\|Page Splits per second | Microsoft SQL Server | False |
| Performance\|Broker Activation\|Stored Procedures Invoked per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Buffer cache hit ratio (%) | Microsoft SQL Server | True |
| Performance\|Buffer Manager\| Checkpoint Pages/sec | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Lazy writes per second | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Page life expectancy | Microsoft SQL Server | True |
| Performance\|Buffer Manager\|Page lookups per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Page reads per second | Microsoft SQL Server | False |
| Performance\|Buffer Manager\|Page writes per second | Microsoft SQL Server | False |
| Performance\|Databases\|Active Transactions | Microsoft SQL Server | True |
| Performance\|Databases\|Data File(s) Size | Microsoft SQL Server | True |
| Performance\|Databases\|Log Bytes Flushed/Sec | Microsoft SQL Server | False |
| Performance\|Databases\|Log File(s) Size | Microsoft SQL Server | False |
| Performance\|Databases\|Log File(s) Used Size | Microsoft SQL Server | False |
| Performance\|Databases\|Log Flush Wait Time | Microsoft SQL Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Performance\|Databases\|Log Flushes per second | Microsoft SQL Server | False |
| Performance\|Databases\| Transactions per second | Microsoft SQL Server | False |
| Performance\|Databases\|Write Transactions per second | Microsoft SQL Server | False |
| Performance\|Databases\|XTP Memory Used | Microsoft SQL Server | False |
| Performance\|General Statistics\|Active temp Tables | Microsoft SQL Server | False |
| Performance\|General Statistics\| Logins per second | Microsoft SQL Server | False |
| Performance\|General Statistics\| Logouts per second | Microsoft SQL Server | False |
| Performance\|General Statistics\| Processes Blocked | Microsoft SQL Server | False |
| Performance\|General Statistics\|Temp Tables Creation Rate | Microsoft SQL Server | False |
| Performance\|General Statistics\|User Connections | Microsoft SQL Server | False |
| Performance\|Locks\|Average Wait Time | Microsoft SQL Server | False |
| Performance\|Locks\|Lock Requests per second | Microsoft SQL Server | False |
| Performance\|Locks\|Lock Wait Time | Microsoft SQL Server | True |
| Performance\|Locks\|Lock Waits per second | Microsoft SQL Server | True |
| Performance\|Locks\|Number of Deadlocks per second | Microsoft SQL Server | True |
| Performance\|Memory Manager\| Connection Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\|Lock Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\|Log Pool Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\| Memory Grants Pending | Microsoft SQL Server | True |
| Performance\|Memory Manager\|SQL Cache Memory | Microsoft SQL Server | False |
| Performance\|Memory Manager\| Target Server Memory | Microsoft SQL Server | True |
| Performance\|Memory Manager\|Total Server Memory | Microsoft SQL Server | True |
| Performance\|Resource Pool Stats\| internal\|Active memory grant amount | Microsoft SQL Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Performance\|Resource Pool Stats\| internal\|CPU Usage Percentage (%) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read Bytes per second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read IO | Microsoft SQL Server | False |
| Wait Stats:<InstanceName>\|Wait Time (ms) | Microsoft SQL Server | False |
| Wait Stats<InstanceName>\|Number of Waiting tasks (ms) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Read IO Throttled Per Second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Write Bytes per second (Bps) | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Disk Write IO Throttled per second | Microsoft SQL Server | False |
| Performance\|Resource Pool Stats\| internal\|Used Memory | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| Batch Requests Per Second | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| SQL Compilations per second | Microsoft SQL Server | False |
| Performance\|SQL Statistics \| SQL Re-Compilations per second | Microsoft SQL Server | False |
| Performance\|Transactions \| Free space in tempdb (KB) | Microsoft SQL Server | False |
| Performance\|Transactions \| Transactions | Microsoft SQL Server | False |
| Performance\|Transactions \| Version Store Size (KB) | Microsoft SQL Server | False |
| Performance\|User Settable Counter \| User Counter 0 to 10 | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Active Requests | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Blocked Tasks | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|CpU Usage (%) | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Queued Requests | Microsoft SQL Server | False |
| Performance\|Workload Group Stats\| internal\|Request Completed/sec | Microsoft SQL Server | False |
| Application Availability | Microsoft SQL Server | False |

There are no metrics collected for Microsoft SQL Server Database.

### MySQL Metrics

Metrics are collected for the MySQL application service.

**Table 433: MySQL Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Aborted connection count | MySQL | True |
| Connection count | MySQL | True |
| Event wait average time | MySQL | False |
| Event wait count | MySQL | False |
| Binary Files\|Binary Files Count | MySQL | False |
| Binary Files\|Binary Size Bytes | MySQL | False |
| Global Status\|Aborted Clients | MySQL | False |
| Global Status\|Binlog Cache Disk Use | MySQL | False |
| Global Status\|Bytes Received | MySQL | False |
| Global Status\|Bytes Sent | MySQL | False |
| Global Status\|Connection Errors Accept | MySQL | False |
| Global Status\|Connection Errors Internal | MySQL | False |
| Global Status\|Connection Errors Max Connections | MySQL | False |
| Global Status\|Queries | MySQL | False |
| Global Status\|Threads Cached | MySQL | False |
| Global Status\|Threads Connected | MySQL | False |
| Global Status\|Threads Running | MySQL | False |
| Global Status\|Uptime | MySQL | False |
| Global Variables\|Delayed Insert Limit | MySQL | False |
| Global Variables\|Delayed Insert Timeout | MySQL | False |
| Global Variables\|Delayed Queue Size | MySQL | False |
| Global Variables\|Max Connect Errors | MySQL | False |
| Global Variables\|Max Connections | MySQL | False |
| Global Variables\|Max Delayed Threads | MySQL | False |
| Global Variables\|Max Error Count | MySQL | False |
| InnoDB\|All deadlock count | MySQL | False |
| InnoDB\|Buffer Pool Bytes Data | MySQL | False |
| InnoDB\|Buffer Pool Bytes Data | MySQL | False |
| InnoDB\|Buffer Pool Bytes Dirty | MySQL | False |
| InnoDB\|Buffer Pool Dump Status | MySQL | False |
| InnoDB\|Buffer Pool Load Status | MySQL | False |
| InnoDB\|Buffer Pool Pages Data | MySQL | False |
| InnoDB\|Buffer Pool Pages Dirty | MySQL | False |
| InnoDB\|Buffer Pool Pages Flushed | MySQL | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| InnoDB\|Buffer pool size | MySQL | True |
| InnoDB\|Checksums | MySQL | False |
| InnoDB\|Open file count | MySQL | False |
| InnoDB\|Row lock average time | MySQL | False |
| InnoDB\|Row lock current waits | MySQL | False |
| InnoDB\|Row lock maximum time | MySQL | False |
| InnoDB\|Row lock time | MySQL | False |
| InnoDB\|Row lock waits | MySQL | True |
| InnoDB\|Table lock count | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Delete | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Fetch | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Insert | MySQL | False |
| Performance Table IO Waits\|IO Waits Total Update | MySQL | False |
| Process List\|Connections | MySQL | False |
| Application Availability | MySQL | False |
| IO waits average time | MySQL Database | False |
| IO waits count | MySQL Database | True |
| Read high priority average time | MySQL Database | False |
| Read high priority count | MySQL Database | False |
| Write concurrent insert average time | MySQL Database | False |
| Write concurrent insert count | MySQL Database | False |

### NGINX Metrics

Metrics are collected for the NGINX application service.

**Table 434: NGINX Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| HTTP Status Info\|Accepts | Nginx | True |
| HTTP Status Info\|Active connections | Nginx | False |
| HTTP Status Info\|Handled | Nginx | True |
| HTTP Status Info\|Reading | Nginx | False |
| HTTP Status Info\|Requests | Nginx | False |
| HTTP Status Info\|Waiting | Nginx | True |
| HTTP Status Info\|Writing | Nginx | False |
| Application Availability | Nginx | False |

### Network Time Protocol Metrics

Metrics are collected for the Network Time Protocol application service.

**Table 435: Network Time Protocol Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| NTPD \|delay | Network Time Protocol | True |
| NTPD \| jitter | Network Time Protocol | True |
| NTPD \| offset | Network Time Protocol | True |
| NTPD \| poll | Network Time Protocol | False |
| NTPD \| reach | Network Time Protocol | True |
| NTPD \| when | Network Time Protocol | False |
| Application Availability | Network Time Protocol | False |

### Oracle WebLogic Server Metrics

Metrics are collected for the Oracle WebLogic Server application service.

**Table 436: Oracle WebLogic Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|Process Cpu Load | Oracle WebLogic Server | True |
| UTILIZATION\|System Cpu Load | Oracle WebLogic Server | False |
| UTILIZATION\|System Load Average | Oracle WebLogic Server | False |
| Application Availability | Oracle WebLogic Server | False |
| UTILIZATION\|Collection Time | Weblogic Garbage Collector | True |
| UTILIZATION\|Connections HighCount | Weblogic JMS Runtime | True |
| UTILIZATION\|JMS Servers TotalCount | Weblogic JMS Runtime | False |
| UTILIZATION\|Active Total Count Used | Weblogic JTA Runtime | False |
| UTILIZATION\|Active Transactions TotalCount | Weblogic JTA Runtime | False |
| UTILIZATION\|Transaction Abandoned TotalCount | Weblogic JTA Runtime | True |
| UTILIZATION\|Transaction RolledBack App TotalCount | Weblogic JTA Runtime | True |
| UTILIZATION\|Heap Memory Usage | Weblogic JVM Memory | True |
| UTILIZATION\|Non Heap Memory Usage | Weblogic JVM Memory | False |
| UTILIZATION\|Peak Usage | Weblogic JVM Memory Pool | True |
| UTILIZATION\|Usage | Weblogic JVM Memory Pool | False |
| UTILIZATION\|UpTime | Weblogic JVM Runtime | False |

### Pivotal TC Server Metrics

Metrics are collected for the Pivotal TC Server application service.

**Table 437: Pivotal TC Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Buffer Pool<InstanceName>\|Count | Pivotal TC Server | False |
| Buffer Pool<InstanceName>\|Memory Used | Pivotal TC Server | False |
| Buffer Pool<InstanceName>\|Total Capacity | Pivotal TC Server | False |
| Class Loading\|Loaded Class Count | Pivotal TC Server | False |
| Class Loading\|Total Loaded Class Count | Pivotal TC Server | False |
| Class Loading\|Unloaded Class Count | Pivotal TC Server | False |
| File Descriptor Usage\|Max File Descriptor Count | Pivotal TC Server | False |
| File Descriptor Usage\|Open File Descriptor Count | Pivotal TC Server | False |
| Garbage Collection:<InstanceName>\| Total Collection Count | Pivotal TC Server | False |
| Garbage Collection:<InstanceName>\|Total Collection Time | Pivotal TC Server | False |
| Process CPU Usage (%) | Pivotal TC Server | True |
| JVM Memory\|Heap Memory Usage\| Committed Memory | Pivotal TC Server | True |
| JVM Memory\|Heap Memory Usage\| Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Heap Memory Usage\| Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Heap Memory Usage\| Used Memory | Pivotal TC Server | True |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | Pivotal TC Server | True |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | Pivotal TC Server | True |
| JVM Memory\|Number of Object Pending Finalization Count | Pivotal TC Server | True |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Committed Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Initial Memory | Pivotal TC Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Peak Usage\|Used Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Committed Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Initial Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Maximum Memory | Pivotal TC Server | False |
| JVM Memory\|Pool:<InstanceName>\| Usage\|Used Memory | Pivotal TC Server | False |
| Process CPU Usage (%) | Pivotal TC Server | True |
| System CPU Usage (%) | Pivotal TC Server | True |
| Uptime | Pivotal TC Server | True |
| Threading\|Thread Count | Pivotal TC Server | False |
| System Load Average | Pivotal TC Server | False |
| Application Availability | Pivotal TC Server | False |
| Current Thread Count | Pivotal TC Server Thread Pool | False |
| Current Threads Busy | Pivotal TC Server Thread Pool | True |
| Total Request Bytes Received | Pivotal TC Server Thread Pool | False |
| Total Request Bytes Sent | Pivotal TC Server Thread Pool | False |
| Total Request Count | Pivotal TC Server Thread Pool | True |
| Total Request Error Count | Pivotal TC Server Thread Pool | True |
| Total Request Processing Time | Pivotal TC Server Thread Pool | True |
| JSP Count | Pivotal TC Server Web Module | False |
| JSP Reload Count | Pivotal TC Server Web Module | False |
| JSP Unload Count | Pivotal TC Server Web Module | False |

### PostgreSQL

Metrics are collected for the PostgreSQL application service.

**Table 438: PostgreSQL**

| Metric Name | Category | KPI |
|---|---|---|
| Buffers\|Buffers Allocated | PostgreSQL | False |
| Buffers\|Buffers Written by Backend | PostgreSQL | True |
| Buffers\|Buffers Written by Background Writer | PostgreSQL | True |
| Buffers\|Buffers Written During Checkpoints | PostgreSQL | True |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Buffers\|fsync Call Executed by Backend | PostgreSQL | False |
| Checkpoints\|Checkpoints sync time | PostgreSQL | False |
| Checkpoints\|Checkpoints write time | PostgreSQL | False |
| Checkpoints\|Requested checkpoints performed count | PostgreSQL | False |
| Checkpoints\|Scheduled checkpoints performed count | PostgreSQL | False |
| Clean scan stopped count | PostgreSQL | False |
| Application Availability | PostgreSQL | False |
| Disk Blocks\|Blocks Cache Hits | PostgreSQL Database | False |
| Disk Blocks\|Blocks Read | PostgreSQL Database | False |
| Disk Blocks\|Blocks Read Time | PostgreSQL Database | False |
| Disk Blocks\|Blocks Write Time | PostgreSQL Database | False |
| Statistics\|Backends Connected | PostgreSQL Database | False |
| Statistics\|Data Written by Queries | PostgreSQL Database | True |
| Statistics\|Deadlocks Detected | PostgreSQL Database | True |
| Statistics\|Queries Cancelled | PostgreSQL Database | True |
| Statistics\|Temp Files Created by Queries | PostgreSQL Database | False |
| Transactions\|Transactions Committed | PostgreSQL Database | True |
| Transactions\|Transactions Rolled Back | PostgreSQL Database | True |
| Tuples\|Tuples Deleted | PostgreSQL Database | True |
| Tuples\|Tuples Fetched | PostgreSQL Database | True |
| Tuples\|Tuples Inserted | PostgreSQL Database | True |
| Tuples\|Tuples Returned | PostgreSQL Database | True |
| Tuples\|Tuples Updated | PostgreSQL Database | True |

### RabbitMQ Metrics

Metrics are collected for the RabbitMQ application service.

**Table 439: RabbitMQ Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| CPU\|Limit | RabbitMQ | False |
| CPU\|Used | RabbitMQ | True |
| Disk\|Free | RabbitMQ | False |
| Disk\|Free limit | RabbitMQ | False |
| FileDescriptor\|Total | RabbitMQ | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| FileDescriptor\|Used | RabbitMQ | False |
| Memory\|Limit | RabbitMQ | False |
| Memory\|Used | RabbitMQ | True |
| Messages\|Acked | RabbitMQ | False |
| Messages\|Delivered | RabbitMQ | False |
| Messages\|Delivered get | RabbitMQ | False |
| Messages\|Published | RabbitMQ | False |
| Messages\|Ready | RabbitMQ | False |
| Messages\|Unacked | RabbitMQ | False |
| Socket\|Limit | RabbitMQ | False |
| Socket\|Used | RabbitMQ | True |
| UTILIZATION\|Channels | RabbitMQ | True |
| UTILIZATION\|Connections | RabbitMQ | True |
| UTILIZATION\|Consumers | RabbitMQ | True |
| UTILIZATION\|Exchanges | RabbitMQ | True |
| UTILIZATION\|Messages | RabbitMQ | True |
| UTILIZATION\|Queues | RabbitMQ | True |
| Application Availability | RabbitMQ | False |
| Messages\|Publish in | RabbitMQ Exchange | False |
| Messages\|Publish out | RabbitMQ Exchange | False |
| Consumer Utilisation | RabbitMQ Queue | False |
| Consumers | RabbitMQ Queue | False |
| Memory | RabbitMQ Queue | False |
| Messages\|Ack | RabbitMQ Queue | False |
| Messages\|Ack rate | RabbitMQ Queue | False |
| Messages\|Deliver | RabbitMQ Queue | False |
| Messages\|Deliver get | RabbitMQ Queue | False |
| Messages\|Persist | RabbitMQ Queue | False |
| Messages\|Publish | RabbitMQ Queue | False |
| Messages\|Publish rate | RabbitMQ Queue | False |
| Messages\|Ram | RabbitMQ Queue | False |
| Messages\|Ready | RabbitMQ Queue | False |
| Messages\|Redeliver | RabbitMQ Queue | False |
| Messages\|Redeliver rate | RabbitMQ Queue | False |
| Messages\|Space | RabbitMQ Queue | False |
| Messages\|Unack | RabbitMQ Queue | False |
| Messages\|Unacked | RabbitMQ Queue | False |
| Messages | RabbitMQ Queue | False |

There are no metrics collected for RabbitMQ Virtual Host.

### Riak KV Metrics

Metrics are collected for the Riak KV application service.

**Table 440: Riak KV Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| UTILIZATION\|CPU Average | Riak KV | False |
| UTILIZATION\|Memory Processes | Riak KV | False |
| UTILIZATION\|Memory Total | Riak KV | False |
| UTILIZATION\|Node GETs | Riak KV | True |
| UTILIZATION\|Node GETs Total | Riak KV | False |
| UTILIZATION\|Node PUTs | Riak KV | True |
| UTILIZATION\|Node PUTs Total | Riak KV | False |
| UTILIZATION\|PBC Active | Riak KV | True |
| UTILIZATION\|PBC Connects | Riak KV | True |
| UTILIZATION\|Read Repairs | Riak KV | True |
| UTILIZATION\|vNODE Index Reads | Riak KV | True |
| UTILIZATION\|vNODE Index Writes | Riak KV | True |
| Application Availability | Riak KV | False |

### SharePoint Metrics

Metrics are collected for the SharePoint Server application service.

**Table 441: SharePoint Server Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Sharepoint Foundation\|Active Threads | SharePoint Server | True |
| Sharepoint Foundation\|Current Page Requests | SharePoint Server | False |
| Sharepoint Foundation\|Executing SQL Queries | SharePoint Server | False |
| Sharepoint Foundation\|Executing Time/Page Request | SharePoint Server | True |
| Sharepoint Foundation\|Incoming Page Requests Rate | SharePoint Server | False |
| Sharepoint Foundation\|Object Cache Hit Count | SharePoint Server | False |
| Sharepoint Foundation\|Reject Page Requests Rate | SharePoint Server | False |
| Sharepoint Foundation\|Responded Page Requests Rate | SharePoint Server | True |
| SQL query executing time | SharePoint Server | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| Application Availability | SharePoint Server | False |
| Network\|Received Data Rate | SharePoint Web Server | True |
| Network\|Sent Data Rate | SharePoint Web Server | True |
| Process\|Processor Time (%) | SharePoint Windows Service | False |
| Process\|Threads | SharePoint Windows Service | False |

### WebSphere Metrics

Metrics are collected for the WebSphere application service.

**Table 442: WebSphere Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| Thread Pool\|Active Count\| Current | Thread Pool | False |
| Thread Pool\|Active Count\| High | Thread Pool | False |
| Thread Pool\|Active Count\| Low | Thread Pool | False |
| Thread Pool\|Active Count\| Lower | Thread Pool | False |
| Thread Pool\|Active Count\| Upper | Thread Pool | False |
| JDBC\|Close Count | JDBC | False |
| JDBC\|Create Count | JDBC | False |
| JDBC\|JDBC Pool Size\| Average | JDBC | False |
| JDBC\|JDBC Pool Size\| Current | JDBC | False |
| JDBC\|JDBC Pool Size\| Lower | JDBC | False |
| JDBC\|JDBC Pool Size\| Upper | JDBC | False |
| Garbage Collection<InstanceName>\| Total Collection Count | WebSphere | False |
| Garbage Collection<InstanceName>\| Total Collection Time | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Initial Memory | WebSphere | False |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Category | KPI |
|---|---|---|
| JVM Memory\|Heap Memory Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Heap Memory Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Committed Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Initial Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\|Non Heap Memory Usage\|Used Memory | WebSphere | False |
| JVM Memory\|Number of Object Pending Finalization Count | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Committed Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Initial Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\|Peak Usage\|Used Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\| Usage\|Committed Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\| Usage\|Initial Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\| Usage\|Maximum Memory | WebSphere | False |
| JVM Memory\| Pool<InstanceName>\| Usage\|Used Memory | WebSphere | False |
| Process Cpu Load | WebSphere | False |
| System Cpu Load | WebSphere | False |
| System Load Average | WebSphere | False |
| Application Availability | WebSphere | False |

### Windows Service Metrics

Metrics are collected for Windows services.

**Table 443: Windows Service Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| AVAILABILITY\|Resource Availability | Services | False |
| UTILIZATION\|Memory Usage(%) | Services | False |
| UTILIZATION\|CPU Usage(%) | Services | False |

### Linux Process Metrics

Metrics are collected for Linux services.

**Table 444: Linux Process Metrics**

| Metric Name | Category | KPI |
|---|---|---|
| AVAILABILITY\|Resource Availability | Processes | False |
| UTILIZATION\|Memory Usage (%) | Processes | False |
| UTILIZATION\|CPU Usage (%) | Processes | False |
| UTILIZATION\|Number of Processes | Processes | False |

### Remote Check Metrics

Metrics are collected for object types such as HTTP, ICMP, TCP, and UDP.

### HTTP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for HTTP remote checks.

HTTP Metrics

**Table 445: HTTP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Content Length | False |
| Response Code | False |
| Response Time | True |
| Result Code | False |

### ICMP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for the ICMP object type.

**Table 446: ICMP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Average Response Time | True |
| Packet Loss (%) | False |
| Packets Received | False |
| Packets Transmitted | False |
| Result Code | False |

### TCP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for the TCP object type.

**Table 447: TCP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Response Time | True |
| Result Code | False |

### UDP Metrics

VMware Aria OperationsVMware Cloud Foundation Operations discovers metrics for the UDP object type.

**Table 448: UDP Metrics**

| Metric Name | KPI |
|---|---|
| Availability | False |
| Response Time | True |
| Result Code | False |

### VeloCloud Application Service Metrics

Metrics are collected for application services supported by VeloCloud.

### VeloCloud Gateway Metrics

Metrics are collected for the VeloCloud Gateway.

**Table 449: VeloCloud Gateway Metrics**

| Component | Metrics |
|---|---|
| DPDK | DPDK:mbuf \| pool free |
| NAT | NAT \| Active Flows (%) |
| | NAT \| Active Flows |
| | NAT \| Active Routes |
| | NAT \| Active Routes Used (%) |
| | NAT \| Connected Peers |
| | NAT \| NAT Entries |
| NTP Server | NTP Server:ntp.ubuntu.com \| offset value |
| Summary | Summary \| Active Tunnels Count (%) |
| | Summary \| Average Packets Dropped |
| | Summary \| Average wMarkDrop |
| | Summary \| BGP Enabled VRFs |
| | Summary \| BGP Neighbors |
| | Summary \| CLR Count |
| | Summary \| Connected Edges |
| | Summary \| NAT |
| | Summary \| SSH Failed Login |
| | Summary \| Unstable Path Percentage |
| | Summary \| VMCP CTRL Drop Count |
| | Summary \| VMCP TX Drop Count |
| VC Queue | VC Queue \| ipv4_bh packet drop |
| VCMP Tunnel | VCMP Tunnel \| ctrl_0 packet drop |
| | VCMP Tunnel \| ctrl_1 packet drop |
| | VCMP Tunnel \| data_0 packet drop |
| | VCMP Tunnel \| data_1 packet drop |
| | VCMP Tunnel \| init packet drop |

### VeloCloud Orchestrator Metrics

Metrics are collected for the VeloCloud Orchestrator.

**Table 450: VeloCloud Orchestrator Metrics**

| Component | Metrics |
|---|---|
| General | General \| Free Memory (%) |
| | General \| Status |

### Metrics - Ngnix

Metrics are collected for the VeloCloud Ngnix.

**Table 451: Ngnix Metrics**

| Component | Metrics |
|---|---|
| HTTP Status Info | HTTP Status Info \| Accepts |
| | HTTP Status Info \| Active Connections |
| | HTTP Status Info \| Handled |
| | HTTP Status Info \| Reading |
| | HTTP Status Info \| Requests |
| | HTTP Status Info \| Waiting |
| | HTTP Status Info \| Writing |

### Metrics - Redis

Metrics are collected for the VeloCloud Redis.

**Table 452: Redis Metrics**

| Component | Metrics |
|---|---|
| Publish Subscribe. | Publish Subscribe \| Channels |
| Total | Total \| Commands Processed |
| | Total \| Connections Received |
| Used | Used \| CPU |
| | Used \| Memory |
| | Used \| Peak Memory |

### Metrics - ClickHouse

Metrics are collected for the VeloCloud Clickhouse.

**Table 453: Clickhouse Metrics**

| Component | Metrics |
|---|---|
| Background | Background \| Pool Task |
| Buffer | Buffers \| Allocation (Bytes) |
| | Buffers \| Compressed Read Buffer (Bytes) |
| | Buffers \| Compressed Read Buffer Blocks |
| | Buffers \| IO Allocation (Bytes) |
| | Buffers \| Storage Buffer (Bytes) |
| | Buffers \| Storage Buffer Rows |
| Events | Events \| Context Lock |
| | Events \| Disk Write Elapsed (µs) |
| | Events \| File Open |
| | Events \| Function Execute |

*Table continued on next page*

*Continued from previous page*

| Component | Metrics |
|---|---|
| | Events \| Hard Page Faults |
| | Events \| Lock Readers Wait (µs) |
| | Events \| OS IO wait (ms) |
| | Events \| OS Write (Bytes) |
| | Events \| Query |
| | Events \| Readers Wait (ms) |
| | Events \| Real Time |
| | Events \| Soft Page Faults (µs) |
| | Events \| System Time (µs) |
| | Events \| User Time (µs) |
| Global Thread | Global \| Global Thread |
| | Global \| Global Thread Active |
| Local Thread | Local \| Local Thread |
| | Local \| Local Thread Active |
| Replicas | Replicas \| Max Absolute Delay |
| | Replicas \| Max Insert In Queue |
| | Replicas \| Max Merge In Queue |
| | Replicas \| Max Queue Size |
| | Replicas \| Max Relative Delay |
| | Replicas \| Total Insert In Queue |
| | Replicas \| Total Merge Queues |
| | Replicas \| Total Queue Size |
| Summary | Summary \| Background Pool Task |
| | Summary \| Dict Cache Requests |
| | Summary \| File Open Writes |
| | Summary \| Merge |
| | Summary \| Number of Databases |
| | Summary \| Number of Distributed Send |
| | Summary \| Number of Tables |
| | Summary \| Read |
| | Summary \| Replicated Checks |
| | Summary \| Storage Buffer Rows |
| | Summary \| Uncompressed Cache Cells |
| | Summary \| Uptime |
| | Summary \| Write |
| | Summary \| Zookeeper Session |
| | Summary \| Zookeeper Watch |
| Write Buffer | Write Buffer \| File Descriptor Write |
| Replicated | Replicated Fetch |
| Memory | Memory Tracking |
| Query | Query Thread |

### Service Discovery Metrics

Service discovery discovers metrics for several objects. It also discovers CPU and memory metrics for discovered services.

### Virtual Machine Metrics

Service Discovery discovers metrics for virtual machines.

**Table 454: Virtual Machine Metrics**

| Metric Name | Description |
|---|---|
| Guest OS Services\|Total Number of Services | Number of out-of-the-box and user-defined services discovered in the VM. |
| Guest OS Services\|Number of User Defined Services | Number of user-defined services discovered in the VM. |
| Guest OS Services\|Number of OOTB Services | Number of out-of-the-box services discovered in the VM. |
| Guest OS Services\|Number of Outgoing Connections | Number of outgoing connection counts from the discovered services. |
| Guest OS Services\|Number of Incoming Connections | Number of incoming connection counts to the discovered services. |

### Service Summary Metrics

Service discovery discovers summary metrics for the service object. The object is a single service object.

**Table 455: Service Summary Metrics**

| Metric Name | Description |
|---|---|
| Summary\|Incoming Connections Count | Number of incoming connections. |
| Summary\|Outgoing Connections Count | Number of outgoing connections. |
| Summary\|Connections Count | Number of incoming and outgoing connections. |
| Summary\|Pid | Process ID. |

### Service Performance Metrics

Service discovery discovers performance metrics for the service object. The object is a single service object.

**Table 456: Service Performance Metrics**

| Metric Name | Description |
|---|---|
| Performance metrics group\|CPU | CPU usage in percentage. |
| Performance metrics group\|Memory | Memory usage in KB. |
| Performance metrics group\|IO Read Throughput | IO read throughput in KBps. |
| Performance metrics group\|IO Write Throughput | IO write throughput in KBps. |

### Service Type Metrics

Service discovery discovers metrics for service type objects.

**Table 457: Service Type Metrics**

| Metric Name | Description |
|---|---|
| Number of instances | Number of instances of this service type. |

### Calculated Metrics

VMware Aria OperationsVMware Cloud Foundation Operations calculates metrics for capacity, badges, and the health of the system. Calculated metrics apply to a subset of objects found in the `describe.xml` file that describes each adapter.

From data that the vCenter adapter collects, VMware Aria OperationsVMware Cloud Foundation Operations calculates metrics for objects of type:

- vSphere World
- Virtual Machine
- Host System
- Datastore

From data that the VMware Aria OperationsVMware Cloud Foundation Operations adapter collects, VMware Aria OperationsVMware Cloud Foundation Operations calculates metrics for objects of type:

- Node
- Cluster

### Capacity Analytics Generated Metrics

The capacity engine computes and publishes metrics that can be found in the Capacity Analytics Generated group. These metrics help you to plan your resource use based on consumer demand.

#### Capacity Analytics Generated Metrics Group

Capacity analytics uses the capacity engine to analyze historical utilization and generate projected utilization. The engine takes the Demand and Usable Capacity (Total Capacity - HA - buffer) metrics as input and calculates the output metrics that belong to the capacity analytics generated metrics group.

The capacity analytics generated metrics group contains containers and each container contains three output metrics, which are Capacity Remaining, Recommended Size, and Recommended Total Capacity. It also contains the Capacity Remaining Percentage and Time Remaining metrics, which show the most constrained values of the containers.

For the capacity metrics group, full metric names include the name of the resource container. For example, if recommended size metrics are computed for CPU or memory, the actual metric names appear as cpu|demand|recommendedSize or mem|demand|recommendedSize.

**Table 458: Capacity Metrics Group**

| Metric Name | Description |
|---|---|
| Time Remaining (Day(s)) | The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: timeRemaining |
| Capacity Remaining | Capacity remaining is the maximum point between the usable capacity now and the projected utilization for 3 days into the future. If the projected utilization is above 100% of the usable capacity, Capacity Remaining is 0.<br>Key: capacityRemaining |
| Capacity Remaining Percentage (%) | The percentage of Capacity Remaining of the most constrained resource with respect to the usable capacity.<br>Key: capacityRemainingPercentage |
| Recommended Size | The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. The warning threshold is the period during which the time remaining is green. Recommended Size excludes HA settings.<br>Key: recommendedSize |
| Recommended Total Capacity | The maximum projected utilization for the projection period from the current time to 30 days after the warning threshold value for time remaining. Recommended Total Capacity excludes HA settings.<br>Key: recommendedTotalCapacity |

## Capacity Analytics Generated Allocation Metrics

Capacity allocation metrics provide information about the allotment of capacity for Cluster Compute and Datastore Cluster Resources.

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|CPU\|Allocation\|Capacity Remaining (vCPUs) | For vSphere objects published on Cluster Compute Resource only. Capacity Remaining based on overcommit ratio (if configured in effective policy).<br>Key: OnlineCapacityAnalytics\|cpu\|alloc\|capacityRemaining |
| Capacity Analytics Generated\|CPU\|Allocation\| Recommended Total Capacity (Cores) | For vSphere objects published on Cluster Compute Resource only. The recommended level of total capacity, to maintain a green state for time remaining for the given object.<br>Key: OnlineCapacityAnalytics\|cpu\|alloc\| recommendedTotalSize |
| Capacity Analytics Generated\|CPU\|Allocation\|Time Remaining (Day(s)) | For vSphere objects published on Cluster Compute Resource only. The number of days remaining is calculated for both group and container. It calculates the time remaining before the resources run out.<br>Key: OnlineCapacityAnalytics\|cpu\|alloc\|timeRemaining |
| CPU\|Allocation\|Usable Capacity after HA and Buffer (vCPUs) | For vSphere objects published on Cluster Compute Resource only. The usable capacity (total capacity - HA) based on configured overcommit ratio.<br>Key: cpu\|alloc\|usableCapacity |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|CPU\|Allocation\| Recommended Size (Cores) | For vSphere objects published on Cluster Compute Resource only. The recommended level of usable capacity (total capacity - HA), to maintain a green state for time remaining for the given object.<br>Key: OnlineCapacityAnalytics\|cpu\|alloc\|recommendedSize |
| vRealize Operations Manager Generated Properties\|CPU\| Allocation\|Overcommit Ratio Setting | For vSphere objects published on Cluster Compute Resource only. This property shows the allocation overcommit ratio for CPU provided in effective policy.<br>Key: System Properties\|cpu\|alloc\|overcommitRatioSetting |
| vRealize Operations Manager Generated Properties\|CPU\| Allocation\|Buffer (%) | CPU buffer percent defined by policy setting for allocation based capacity computation.<br>Key: Properties\|cpu\|alloc\|bufferSetting |
| Capacity Analytics Generated\|Memory\|Allocation\|Capacity Remaining (KB) | For vSphere objects published on Cluster Compute Resource only. Capacity Remaining based on overcommit ratio (if configured in effective policy).<br>Key: OnlineCapacityAnalytics\|mem\|alloc\| capacityRemaining |
| Capacity Analytics Generated\|Memory\|Allocation\| Recommended Total Capacity (KB) | For vSphere objects published on Cluster Compute Resource only. The recommended level of total capacity, to maintain a green state for time remaining for the given object.<br>Key: OnlineCapacityAnalytics\|mem\|alloc\| recommendedTotalSize |
| Capacity Analytics Generated\|Memory\|Allocation\|Time Remaining (Day(s)) | For vSphere objects published on Cluster Compute Resource only. The number of days remaining is calculated for both group and container. It calculates the time remaining before the resources run out.<br>Key: OnlineCapacityAnalytics\|mem\|alloc\|timeRemaining |
| Memory\|Allocation\|Usable Capacity (KB) | For vSphere objects published on Cluster Compute Resource only. The usable capacity (total capacity - HA) based on configured overcommit ratio.<br>Key: mem\|alloc\|usableCapacity |
| Capacity Analytics Generated\|Memory\|Allocation\| Recommended Size (KB) | For vSphere objects published on Cluster Compute Resource only. The recommended level of usable capacity (total capacity - HA), to maintain a green state for time remaining for the given object.<br>Key: OnlineCapacityAnalytics\|mem\|alloc\| recommendedSize |
| vRealize Operations Manager Generated Properties\| Memory\|Allocation\|Overcommit Ratio Setting | For vSphere objects published on Cluster Compute Resource only. This property shows the allocation overcommit ratio for Memory provided in effective policy.<br>Key: System Properties\|mem\|alloc\|overcommitRatioSetting |
| vRealize Operations Manager Generated Properties\| Memory\|Allocation\|Buffer (%) | Memory buffer percent defined by policy setting for allocation based capacity computation.<br>Key: System Properties\|mem\|alloc\|bufferSetting |
| Capacity Analytics Generated\|Disk Space\|Allocation\| Capacity Remaining (GB) | For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. Capacity |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Remaining based on overcommit ratio (if configured in effective policy).<br>Key: OnlineCapacityAnalytics\|diskspace\|alloc\|capacityRemaining |
| Capacity Analytics Generated\|Disk Space\|Allocation\|Recommended Size (GB) | For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. The recommended level of total capacity to maintain a green state for time remaining for the given object.<br>Key: OnlineCapacityAnalytics\|diskspace\|alloc\|recommendedSize |
| Capacity Analytics Generated\|Disk Space\|Allocation\|Time Remaining (Day(s)) | For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. The number of days remaining is calculated for both group and container. It calculates the time remaining before the resources run out.<br>Key: OnlineCapacityAnalytics\|diskspace\|alloc\|timeRemaining |
| Disk Space\|Allocation\|Usable Capacity (GB) | For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. Usable capacity based on overcommit ratio (if configured in effective policy).<br>Key: diskspace\|alloc\|usableCapacity |
| vRealize Operations Manager Generated Properties\|Disk Space\|Allocation\|Overcommit Ratio Setting | For vSphere objects published on Cluster Compute Resource and Datastore Cluster Resource. This property shows the allocation overcommit ratio for Disk Space provided in effective policy.<br>key: System Properties\|diskspace\|alloc\|overcommitRatioSetting |
| vRealize Operations Manager Generated Properties\|Disk Space\|Allocation\|Buffer (%) | Disk Space buffer percent defined by policy setting for allocation based capacity computation.<br>Key: System Properties\|diskspace\|alloc\|bufferSetting |

## Capacity Analytics Generated Profiles Metrics

Profiles metrics provide information about the profile specific capacity for Cluster Compute, Datastore Cluster, Data Center, Custom Data Center, and vCenter Server resources.

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|Capacity Remaining (Profile) | Published on Cluster Compute Resource. Calculated as a minimum of all Profiles\|capacityRemainingProfile_<profile uuid> metrics.<br>Key: OnlineCapacityAnalytics\|capacityRemainingProfile |
| Capacity Analytics Generated\|Capacity Remaining (Profile) | Published on Datastore Cluster Resource. Calculated as a minimum of all Profiles\|capacityRemainingProfile_<profile uuid> metrics.<br>Key: OnlineCapacityAnalytics\|capacityRemainingProfile |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|Capacity Remaining (Profile) | Published on Data Center, Custom Data Center and vCenter Server Resources. Computed as a sum of OnlineCapacityAnalytics\|capacityRemainingProfile metric of descendant Cluster Compute Resources.<br>Key: OnlineCapacityAnalytics\|capacityRemainingProfile |

**Capacity Demand Model Metrics**

Demand model metrics provide information about the usable capacity and projected utilization of resources across VMs, Host Systems, Cluster Compute, Datastore Cluster, Data Center, Custom Data Center, and vCenter Server resources.

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|CPU\|Capacity Remaining (MHz) | Published on Virtual Machine. The max point between the usable capacity and the projected utilization between now and three days.<br>Key: OnlineCapacityAnalytics\|cpu\|capacityRemaining |
| Capacity Analytics Generated\|CPU\|Recommended Size (MHz) | Published on Virtual Machine. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|cpu\|recommendedSize |
| Capacity Analytics Generated\|CPU\|Time Remaining (Day(s)) | Published on Virtual Machine. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|cpu\|timeRemaining |
| Capacity Analytics Generated\|Disk Space\|Capacity Remaining (GB) | Published on Virtual Machine. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|diskspace\|capacityRemaining |
| Capacity Analytics Generated\|Disk Space\|Recommended Size (GB) | Published on Virtual Machine. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|diskspace\|recommendedSize |
| Capacity Analytics Generated\|Disk Space\|Time Remaining (Day(s)) | Published on Virtual Machine. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|diskspace\|timeRemaining |
| Capacity Analytics Generated\|Memory\|Capacity Remaining (KB) | Published on Virtual Machine. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|mem\|capacityRemaining |
| Capacity Analytics Generated\|Memory\|Recommended Size (KB) | Published on Virtual Machine. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|mem\|recommendedSize |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|Memory\|Time Remaining (Day(s)) | Published on Virtual Machine. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|mem\|timeRemaining |
| Capacity Analytics Generated\|CPU\|Demand\|Capacity Remaining (MHz) | Published on Host System. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|capacityRemaining |
| vRealize Operations Manager Generated Properties\|CPU\|Demand\|Buffer (%) | CPU buffer percent defined by policy setting for demand based capacity computation.<br>Key: System Properties\|cpu\|demand\|bufferSetting |
| Capacity Analytics Generated\|CPU\|Demand\|Recommended Size (MHz) | Published on Host System. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|recommendedSize |
| Capacity Analytics Generated\|CPU\|Demand\|Time Remaining (Day(s)) | Published on Host System. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|timeRemaining |
| Capacity Analytics Generated\|Disk Space\|Demand\|Capacity Remaining (GB) | Published on Host System. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|capacityRemaining |
| vRealize Operations Manager Generated Properties\|Disk Space\|Demand\|Buffer (%) | Disk Space buffer percent defined by policy setting for demand based capacity computation.<br>System Properties\|diskspace\|demand\|bufferSetting |
| Capacity Analytics Generated\|Disk Space\|Demand\|Recommended Size (GB) | Published on Host System. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|recommendedSize |
| Capacity Analytics Generated\|Disk Space\|Demand\|Time Remaining (Day(s)) | Published on Host System. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|timeRemaining |
| Capacity Analytics Generated\|Memory\|Demand\|Capacity Remaining (KB) | Published on Host System. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|mem\|demand\|capacityRemaining |
| vRealize Operations Manager Generated Properties\|Memory\|Demand\|Buffer (%) | Memory buffer percent defined by policy setting for demand based capacity computation.<br>Key: System Properties\|mem\|demand\|bufferSetting |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|Memory\|Demand\| Recommended Size (KB) | Published on Host System. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|mem\|demand\| recommendedSize |
| Capacity Analytics Generated\|Memory\|Demand\|Time Remaining (Day(s)) | Published on Host System. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|mem\|demand\| timeRemaining |
| Capacity Analytics Generated\|Disk Space\|Usage\|Capacity Remaining (GB) | Published on Datastore. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|diskspace\|total\| capacityRemaining |
| Capacity Analytics Generated\|Disk Space\|Usage\| Recommended Size (GB) | Published on Datastore. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|diskspace\|total\| recommendedSize |
| Capacity Analytics Generated\|Disk Space\|Usage\|Time Remaining (Day(s)) | Published on Datastore. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|diskspace\|total\| timeRemaining |
| Capacity Analytics Generated\|CPU\|Demand\|Capacity Remaining (MHz) | Published on Cluster Compute Resource. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\| capacityRemaining |
| Capacity Analytics Generated\|CPU\|Demand\| Recommended Size (MHz) | Published on Cluster Compute Resource. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the remaining time.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\| recommendedSize |
| Capacity Analytics Generated\|CPU\|Demand\| Recommended Total Capacity (MHz) | Published on Cluster Compute Resource. The recommended level of total capacity to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\| recommendedTotalSize |
| Capacity Analytics Generated\|CPU\|Demand\|Time Remaining (Day(s)) | Published on Cluster Compute Resource. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|timeRemaining |
| Capacity Analytics Generated\|Disk Space\|Demand\| Capacity Remaining (GB) | Published on Cluster Compute Resource. The max point between the usable capacity and the projected utilization between now and three days into the future. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: OnlineCapacityAnalytics\|diskspace\|demand\|capacityRemaining |
| Capacity Analytics Generated\|Disk Space\|Demand\|Recommended Size (GB) | Published on Cluster Compute Resource. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|recommendedSize |
| Capacity Analytics Generated\|Disk Space\|Demand\|Time Remaining (Day(s)) | Published on Cluster Compute Resource. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|timeRemaining |
| Capacity Analytics Generated\|Memory\|Demand\|Capacity Remaining (KB) | Published on Cluster Compute Resource. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|mem\|demand\|capacityRemaining |
| Capacity Analytics Generated\|Memory\|Demand\|Recommended Size (KB) | Published on Cluster Compute Resource. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|mem\|demand\|recommendedSize |
| Capacity Analytics Generated\|Memory\|Demand\|Recommended Total Capacity (KB) | Published on Cluster Compute Resource. The recommended level of total capacity to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|mem\|demand\|recommendedTotalSize |
| Capacity Analytics Generated\|Memory\|Demand\|Time Remaining (Day(s)) | Published on Cluster Compute Resource. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|mem\|demand\|timeRemaining |
| Capacity Analytics Generated\|Disk Space\|Usage\|Capacity Remaining (GB) | Published on Datastore Cluster. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|diskspace\|total\|capacityRemaining |
| Capacity Analytics Generated\|Disk Space\|Usage\|Recommended Size (GB) | Published on Datastore Cluster. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|diskspace\|total\|recommendedSize |
| Capacity Analytics Generated\|Disk Space\|Usage\|Time Remaining (Day(s)) | Published on Datastore Cluster. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|diskspace\|total\|timeRemaining |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|CPU\|Demand\|Capacity Remaining (MHz) | Published on Datacenter, Custom Datacenter, vCenter. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|capacityRemaining |
| Capacity Analytics Generated\|CPU\|Demand\|Recommended Size (MHz) | Published on Datacenter, Custom Datacenter, vCenter. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|recommendedSize |
| Capacity Analytics Generated\|CPU\|Demand\|Recommended Total Capacity (MHz) | Published on Datacenter, Custom Datacenter, vCenter. The recommended level of total capacity to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|recommendedTotalSize |
| Capacity Analytics Generated\|CPU\|Demand\|Time Remaining (Day(s)) | Published on Datacenter, Custom Datacenter, vCenter. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|cpu\|demand\|timeRemaining |
| Capacity Analytics Generated\|Disk Space\|Demand\|Capacity Remaining (GB) | Published on Datacenter, Custom Datacenter, vCenter. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|capacityRemaining |
| Capacity Analytics Generated\|Disk Space\|Demand\|Recommended Size (GB) | Published on Datacenter, Custom Datacenter, vCenter. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|recommendedSize |
| Capacity Analytics Generated\|Disk Space\|Demand\|Time Remaining (Day(s)) | Published on Datacenter, Custom Datacenter, vCenter. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|diskspace\|demand\|timeRemaining |
| Capacity Analytics Generated\|Memory\|Demand\|Capacity Remaining (KB) | Published on Datacenter, Custom Datacenter, vCenter. The max point between the usable capacity and the projected utilization between now and three days into the future.<br>Key: OnlineCapacityAnalytics\|mem\|demand\|capacityRemaining |
| Capacity Analytics Generated\|Memory\|Demand\|Recommended Size (KB) | Published on Datacenter, Custom Datacenter, vCenter. The recommended level of usable capacity (total capacity - HA) to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|mem\|demand\|recommendedSize |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| Capacity Analytics Generated\|Memory\|Demand\| Recommended Total Capacity (KB) | Published on Datacenter, Custom Datacenter, vCenter. The recommended level of total capacity to maintain a green state for the time remaining.<br>Key: OnlineCapacityAnalytics\|mem\|demand\| recommendedTotalSize |
| Capacity Analytics Generated\|Memory\|Demand\|Time Remaining (Day(s)) | Published on Datacenter, Custom Datacenter, vCenter. The number of days remaining till the projected utilization crosses the threshold for the usable capacity.<br>Key: OnlineCapacityAnalytics\|mem\|demand\| timeRemaining |

### Badge Metrics

Badge metrics provide information for badges in the user interface. They report the health, risk, and efficiency of objects in your environment.

VMware Aria OperationsVMware Cloud Foundation Operations 6.x analyzes badge metric data at five-minute averages, instead of hourly. As a result, you might find that efficiency and risk badge calculations are more sensitive than in previous versions. Badge metrics continue to be published nightly.

**Table 459: Badge Metrics**

| Metric Name | Description |
|---|---|
| Badge\|Compliance | Badge\|Compliance(%) metric shows the compliance score for the given object based on the number of violated and total number of compliance symptoms calculated with the following formula `[Math.round(100 - (((double)triggeredSymptoms/ totalSymptoms)*100))]`. So, this metric shows the compliance score per object based on the violated symptoms (which can be seen under the object's details Compliance tab). This metric should not be mixed up with the compliance score shown for the benchmark in the **Operations › Compliance** page which considers compliant/ non-compliant vs total objects.Badge\|Compliance(%) metric is used in 4 views: - Compliance \ VM Distribution - Compliance \ vSphere Distributed Port Groups - Compliance \ vSphere ESXi Hosts - Compliance \ vSphere VMs, as well as 1 deprecated dashboard's widget - vSphere Security Compliance dashboard's Compliance Summary widget. |
| Badge\|Efficiency | Overall score for efficiency. The final score is between 1-100. Where Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1. The score is derived from the criticality of alerts in the Efficiency category. |
| Badge\|Health | Overall score for health. The final score is between 1-100. Where Green - 100, Yellow - 75, Orange - 50, Red - 25, Unknown: -1. The score is derived from the criticality of alerts in the Health category. |
| Badge\|Risk | Overall score for risk. The final score is between 1-100. Where Green - 0, Yellow - 25, Orange - 50, Red - 75, Unknown: -1. The score is derived from the criticality of alerts in the Risk category. |

### System Metrics

System metrics provide information used to monitor the health of the system. They help you to identify problems in your environment.

**Table 460: System Metrics**

| Metric Name | Description |
|---|---|
| vRealize Operations Generated\|Self - Health Score | This metric displays the system health score of self resource. The value ranges from 0 to 100 depending on noise and the number of alarms.<br>Key: System Attributes\|health |
| vRealize Operations Generated\|Self - Metric Count | This metric displays the number of metrics that the adapter generates for the given object. This value does not include the number of metrics generated by VMware Aria OperationsVMware Cloud Foundation Operations, such as, Badge metrics, vRealize Operations Generated metrics and metrics generated by Capacity Engine<br>Key: System Attributes\|all_metrics |
| vRealize Operations Generated\|Total Anomalies | This metric displays the number of active anomalies (symptoms, events, DT violations) on the object and its children.<br>In previous versions of VMware Cloud Foundation Operations, this metric used to be named vRealize Operations Generated\| Self - Total Anomalies.<br><br>Key: System Attributes\|total_alarms |
| vRealize Operations Generated\|Full Set - Metric Count | This metric displays the number of metrics that the adapter of the children of the given object generates.<br>Key: System Attributes\|child_all_metrics |
| vRealize Operations Generated\|Availability | This metric value is computed based on the adapter instance statuses monitoring the resource. Resource availability is displayed as 0-down, 1-Up, -1-Unknown.<br>Key: System Attributes\|availability |
| vRealize Operations Generated\|Alert Count Critical | This metric displays the number of critical alerts on the object and its children.<br>Key: System Attributes\|alert_count_critical |
| vRealize Operations Generated\|Alert Count Immediate | This metric displays the number of immediate alerts on the object and its children.<br>Key: System Attributes\|alert_count_immediate |
| vRealize Operations Generated\|Alert Count Warning | This metric displays the number of active warning alerts on the object and its children.<br>Key: System Attributes\|alert_count_warning |
| vRealize Operations Generated\|Alert Count Info | This metric displays the number of active info alerts on the object and its children.<br>Key: System Attributes\|alert_count_info |
| vRealize Operations Generated\|Total Alert Count | This metric displays the sum of all alert count metrics.<br>In previous versions of VMware Cloud Foundation Operations, this metric was named vRealize Operations Generated\|Full Set - Alert Count. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Description |
|---|---|
| | Key: System Attributes\|total_alert_count |
| vRealize Operations Generated\|Self-Alert Count | This metric displays the number of all alerts on the object.<br>Key: System Attributes\|self_alert_count |

### Logs Generated Metrics

The metrics in the Logs Generated group provide information that you can use to observe or troubleshoot VMware Cloud Foundation Operations for failures and to monitor performance.

When VMware Cloud Foundation Operations is integrated with Logs and metric calculation is enabled, Logs calculates the number of logs corresponding to different queries and sends them as metrics to VMware Cloud Foundation Operations. These metrics are calculated for vCenter objects, host objects, and virtual machine objects. The metrics can be mapped to a VMware Cloud Foundation Operations object based on the Logs field `vmw_vrops_id`, which is constructed based on hostname or source fields.

**Table 461: Log Insight Generated Metrics**

| Metric Name | Description |
|---|---|
| Logs Generated\|Error Count | The number of error logs for the selected object.<br>Key: log_insight_generated\|error_count |
| Logs Generated\|Total Log Count | The total number of logs for the selected object.<br>Key: log_insight_generated\|total_log_count |
| Logs Generated\|Warning Count | The number of warning logs for the selected object.<br>Key: log_insight_generated\|warning_count |

### Self-Monitoring Metrics for VMware Aria OperationsVMware Cloud Foundation Operations

VMware Aria OperationsVMware Cloud Foundation Operations uses the VMware Aria OperationsVMware Cloud Foundation Operations adapter to collect metrics that monitor its own performance. These self-monitoring metrics drive capacity models for VMware Aria OperationsVMware Cloud Foundation Operations objects and are useful for diagnosing problems with VMware Aria OperationsVMware Cloud Foundation Operations.

### Analytics Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations analytics service, including threshold checking metrics.

**Table 462: Analytics Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| ActiveAlarms | Active DT Symptoms | Active DT Symptoms. |
| ActiveAlerts | Active Alerts | Active alerts. |
| PrimaryResourcesCount | Number of primary objects | Number of primary objects |
| LocalResourcesCount | Number of local objects | Number of local objects |
| PrimaryMetricsCount | Number of primary metrics | Number of primary metrics |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| LocalMetricsCount | Number of local metrics | Number of local metrics |
| ReceivedResourceCount | Number of received objects | Number of received objects |
| ReceivedMetricCount | Number of received metrics | Number of received metrics |
| LocalFDSize | Number of forward data entries | Number of locally stored primary and redundant entries in forward data region. |
| LocalPrimaryFDSize | Number of primary forward data entries | Number of locally stored primary entries in forward data region. |
| LocalFDAltSize | Number of alternative forward data entries | Number of locally stored primary and redundant entries in alternative forward data region. |
| LocalPrimaryFDAltSize | Number of alternative primary forward data entries | Number of locally stored primary entries in alternative forward data region. |
| CurrentHeapSize | Current heap size | Current heap size. |
| MaxHeapSize | Max heap size | Max heap size |
| CommittedMemory | Committed memory | Committed memory |
| CPUUsage | CPU usage | CPU usage |
| Threads | Threads | Threads |
| UpStatus | Threads | Threads |

## Overall Threshold Checking Metrics for the Analytics Service

Overall threshold checking captures various metrics for work items used to process incoming observation data. All metrics keys for the overall threshold checking metrics begin with OverallThresholdChecking, as in `OverallThresholdChecking|Count` or `OverallThresholdChecking|CheckThresholdAndHealth|OutcomeObservationsSize|TotalCount`.

**Table 463: Overall Threshold Checking Metrics for the Analytics Service**

| Metric Key | Metric Name | Description |
|---|---|---|
| Count | Count | Count |
| Duration\|TotalDuration | Total | Total length of duration (ms) |
| Duration\|AvgDuration | Average | Average duration (ms) |
| Duration\|MinDuration | Minimum | Minimum duration (ms) |
| Duration\|MaxDuration | Maximum | Maximum duration (ms) |
| IncomingObservationsSize\|TotalCount | Total | Total |
| IncomingObservationsSize\|AvgCount | Average | Average |
| IncomingObservationsSize\|MinCount | Minimal | Minimal |
| IncomingObservationsSize\|MaxCount | Maximal | Maximal |
| CheckThresholdAndHealth\|Count | Count | Count |
| CheckThresholdAndHealth\|Duration\|TotalDuration | Total | Total length of duration (ms) |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| CheckThresholdAndHealth\|Duration\| AvgDuration | Average | Average duration (ms) |
| CheckThresholdAndHealth\|Duration\| MinDuration | Minimum | Minimum duration (ms) |
| CheckThresholdAndHealth\|Duration\| MaxDuration | Maximum | Maximum duration (ms) |
| CheckThresholdAndHealth\| OutcomeObservationsSize\|TotalCount | Total | Total |
| CheckThresholdAndHealth\| OutcomeObservationsSize\|AvgCount | Average | Average |
| CheckThresholdAndHealth\| OutcomeObservationsSize\|MinCount | Minimal | Minimal |
| CheckThresholdAndHealth\| OutcomeObservationsSize\|MaxCount | Maximal | Maximal |
| SuperMetricComputation\|Count | Count | Count |
| SuperMetricComputation\|Duration\| TotalDuration | Total | Total length of duration (ms) |
| SuperMetricComputation\|Duration\| AvgDuration | Average | Average duration (ms) |
| SuperMetricComputation\|Duration\| MinDuration | Minimum | Minimum duration (ms) |
| SuperMetricComputation\|Duration\| MaxDuration | Maximum | Maximum duration (ms) |
| SuperMetricComputation\| SuperMetricsCount\|TotalCount | Total | Total |
| SuperMetricComputation\| SuperMetricsCount \|AvgCount | Average | Average |
| SuperMetricComputation\| SuperMetricsCount \|MinCount | Minimal | Minimal |
| SuperMetricComputation\| SuperMetricsCount \|MaxCount | Maximal | Maximal |
| StoreObservationToFSDB\|Count | Count | Count |
| StoreObservationToFSDB\|Duration\| TotalDuration | Total | Total length of duration (ms) |
| StoreObservationToFSDB\|Duration\| AvgDuration | Average | Average duration (ms) |
| StoreObservationToFSDB\|Duration\| MinDuration | Minimum | Minimum duration (ms) |
| StoreObservationToFSDB\|Duration\| MaxDuration | Maximum | Maximum duration (ms) |
| StoreObservationToFSDB\| StoredObservationsSize\|TotalCount | Total | Total |
| StoreObservationToFSDB\| StoredObservationsSize\|AvgCount | Average | Average |
| StoreObservationToFSDB\| StoredObservationsSize\|MinCount | Minimal | Minimal |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| StoreObservationToFSDB\|StoredObservationsSize\|MaxCount | Maximal | Maximal |
| UpdateResourceCache\|Count | Count | Count |
| UpdateResourceCache\|Duration\|TotalDuration | Total | Total |
| UpdateResourceCache\|Duration\|AvgDuration | Average | Average |
| UpdateResourceCache\|Duration\|MinDuration | Minimum | Minimum |
| UpdateResourceCache\|Duration\|MaxDuration | Maximum | Maximum |
| UpdateResourceCache\|ModifcationEstimateCount\|TotalCount | Total | The number of estimated modifications done during each resource cache object update. |
| UpdateResourceCache\|ModifcationEstimateCount\|AvgCount | Average | Average |
| UpdateResourceCache\|ModifcationEstimateCount\|MinCount | Minimal | Minimal |
| UpdateResourceCache\|ModifcationEstimateCount\|MaxCount | Maximal | Maximal |
| ManageAlerts\|Count | Count | The total number of times the threshold checking work items perform alert updates. |
| ManageAlerts\|Duration\|TotalDuration | Total | The duration for the alert updates operations. |
| ManageAlerts\|Duration\|AvgDuration | Average | Average |
| ManageAlerts\|Duration\|MinDuration | Minimum | Minimum |
| ManageAlerts\|Duration\|MaxDuration | Maximum | Maximum |
| UpdateSymptoms\|Count | Count | The total number of times the threshold checking work items check and build symptoms. |
| UpdateSymptoms\|Duration\|TotalDuration | Total | The duration for the check and build symptoms operation. |
| UpdateSymptoms\|Duration\|AvgDuration | Average | Average |
| UpdateSymptoms\|Duration\|MinDuration | Minimum | Minimum |
| UpdateSymptoms\|Duration\|MaxDuration | Maximum | Maximum |

## Dynamic Threshold Calculation Metrics for the Analytics Service

All metrics keys for the dynamic threshold calculation metrics begin with DtCalculation, as in `DtCalculation|DtDataWrite|WriteOperationCount` or `DtCalculation|DtAnalyze|AnalyzeOperationCount`.

**Table 464: Dynamic Threshold Calculation Metrics for the Analytics Service**

| Metric Key | Metric Name | Description |
|---|---|---|
| DtDataWrite\|WriteOperationCount | Write operation count | Write operation count |
| DtDataWrite\|Duration\|TotalDuration | Total | Total length of duration (ms) |
| DtDataWrite\|Duration\|AvgDuration | Average | Average duration (ms) |
| DtDataWrite\|Duration\|MinDuration | Minimum | Minimum duration (ms) |
| DtDataWrite\|Duration\|MaxDuration | Maximum | Maximum duration (ms) |
| DtDataWrite\|SavedDtObjectCount\|TotalCount | Total | Total |
| DtDataWrite\|SavedDtObjectCount\|AvgCount | Average | Average |
| DtDataWrite\|SavedDtObjectCount\|MinCount | Minimal | Minimal |
| DtDataWrite\|SavedDtObjectCount\|MaxCount | Maximal | Maximal |
| DtAnalyze\|AnalyzeOperationCount | Analyze Operation Count | Analyze Operation Count |
| DtAnalyze\|Duration\|TotalDuration | Total | Total length of duration (ms) |
| DtAnalyze\|Duration\|AvgDuration | Average | Average duration (ms) |
| DtAnalyze\|Duration\|MinDuration | Minimum | Minimum duration (ms) |
| DtAnalyze\|Duration\|MaxDuration | Maximum | Maximum duration (ms) |
| DtAnalyze\|AnalyzedMetricsCount\|TotalCount | Total | Total |
| DtAnalyze\|AnalyzedMetricsCount\|AvgCount | Average | Average |
| DtAnalyze\|AnalyzedMetricsCount\|MinCount | Minimal | Minimal |
| DtAnalyze\|AnalyzedMetricsCount\|MaxCount | Maximal | Maximal |
| DtDataRead\|ReadOperationsCount | Read Operation Count | Read Operation Count |
| DtDataRead\|Duration\|TotalDuration | Total | Total length of duration (ms) |
| DtDataRead\|Duration\|AvgDuration | Average | Average duration (ms) |
| DtDataRead\|Duration\|MinDuration | Minimum | Minimum duration (ms) |
| DtDataRead\|Duration\|MaxDuration | Maximum | Maximum duration (ms) |
| DtDataRead\|ReadDataPointsCount\|TotalCount | Total | Total |
| DtDataRead\|ReadDataPointsCount\|AvgCount | Average | Average |
| DtDataRead\|ReadDataPointsCount\|MinCount | Minimal | Minimal |
| DtDataRead\|ReadDataPointsCount\|MaxCount | Maximal | Maximal |

**Table 465: Function Call Metrics for the Analytics Service**

| Metric Key | Metric Name | Description |
|---|---|---|
| FunctionCalls\|Count | Number of function calls | Number of function calls |
| FunctionCalls\|AvgDuration | Average execution time | Average execution time |
| FunctionCalls\|MaxDuration | Max execution time | Max execution time |

### Collector Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations Collector service objects.

**Table 466: Collector Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| ThreadpoolThreadsCount | Number of pool threads | Number of pool threads. |
| RejectedFDCount | Number of rejected forward data | Number of rejected forward data |
| RejectedFDAltCount | Number of rejected alternative forward data | Number of rejected alternative forward data |
| SentFDCount | Number of sent objects | Number of sent objects |
| SentFDAltCount | Number of alternative sent objects | Number of alternative sent objects |
| CurrentHeapSize | Current heap size (MB) | Current heap size. |
| MaxHeapsize | Max heap size (MB) | Maximum heap size. |
| CommittedMemory | Committed memory (MB) | Amount of committed memory. |
| CPUUsage | CPU usage | CPU usage. |
| Threads | Threads | Number of threads. |
| UpStatus | Up Status | Up Status |

### Controller Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations Controller objects.

**Table 467: Controller Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| RequestedMetricCount | Number of requested metrics | Number of requested metrics |
| ApiCallsCount | Number of API calls | Number of API calls |
| NewDiscoveredResourcesCount | Number of discovered objects | Number of discovered objects |

### FSDB Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations file system database (FSDB) objects.

**Table 468: FSDB Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| StoragePoolElementsCount | Number of storage work items | Number of storage work items |
| FsdbState | Fsdb state | Fsdb state |
| StoredResourcesCount | Number of stored objects | Number of stored objects |
| StoredMetricsCount | Number of stored metrics | Number of stored metrics |

**Table 469: Storage Thread Pool Metrics for FSDB**

| Metric Key | Metric Name | Description |
|---|---|---|
| StoreOperationsCount | Store operations count | Store operations count |
| StorageThreadPool\|Duration\| TotalDuration | Total | Total number of duration (ms) |
| StorageThreadPool\|Duration\|AvgDuration | Average | Average duration (ms) |
| StorageThreadPool\|Duration\|MinDuration | Minimum | Minimum duration (ms) |
| StorageThreadPool\|Duration\|MaxDuration | Maximum | Maximum duration (ms) |
| StorageThreadPool\|SavedMetricsCount\| TotalCount | Total | Total |
| StorageThreadPool\|SavedMetricsCount\| AvgCount | Average | Average |
| StorageThreadPool\|SavedMetricsCount\| MinCount | Minimal | Minimal |
| StorageThreadPool\|SavedMetricsCount\| MaxCount | Maximal | Maximal |

**Product UI Metrics**

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations product user interface objects.

**Table 470: Product UI Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| ActiveSessionsCount | Active sessions | Active sessions |
| CurrentHeapSize | Current heap size | Current heap size. |
| MaxHeapsize | Max heap size | Maximum heap size. |
| CommittedMemory | Committed memory | Amount of committed memory. |
| CPUUsage | CPU usage | Percent CPU use. |
| Threads | Threads | Number of threads. |
| SessionCount | Number of active sessions | Number of active sessions |
| SelfMonitoringQueueSize | Self Monitoring queue size | Self Monitoring queue size |

**Table 471: API Call Metrics for the Product UI**

| Metric Key | Metric Name | Description |
|---|---|---|
| APICalls\|HTTPRequesterRequestCount | HTTPRequester request count | HTTPRequester request count |
| APICalls\| AvgHTTPRequesterRequestTime | HTTPRequester average request time | HTTPRequester average request time (ms) |
| APICalls\|FailedAuthenticationCount | Failed Authentication Count | Failed Authentication Count |
| APICalls\|AvgAlertRequestTime | Average alert request time | Average alert request time (ms) |
| APICalls\|AlertRequestCount | Alert request count | Alert request count |
| APICalls\|AvgMetricPickerRequestTime | Average metric-picker request time | Average metric-picker request time (ms) |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| APICalls\|MetricPickerRequestCount | Metric picker request count | Metric picker request count |
| APICalls\|HeatmapRequestCount | Heatmap request count | Heatmap request count |
| APICalls\|AvgHeatmapRequestTime | Average HeatMap request time | Average HeatMap request time (ms) |
| APICalls\|MashupChartRequestCount | Mashup Chart request count | Mashup Chart request count |
| APICalls\|AvgMashupChartRequestTime | Average Mashup Chart request time | Average Mashup Chart request time (ms) |
| APICalls\|TopNRequestCount | Top N request count | Top N request count |
| APICalls\|AvgTopNRequestTime | Average Top N request time | Average Top N request time (ms) |
| APICalls\|MetricChartRequestCount | Metric Chart request count | Metric Chart request count |
| APICalls\|AvgMetricChartRequestTime | Average MetricChart request time | Average MetricChart request time (ms) |

### Admin UI Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations administration user interface objects.

**Table 472: Admin UI Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| CurrentHeapSize | Current heap size | Current heap size (MB). |
| MaxHeapsize | Max heap size | Maximum heap size (MB). |
| CommittedMemory | Committed memory | Amount of committed memory (MB) . |
| CPUUsage | CPU usage | CPU usage (%). |
| Threads | Threads | Number of threads. |
| SessionCount | Number of active sessions | Number of active sessions |
| SelfMonitoringQueueSize | Self Monitoring queue size | Self Monitoring queue size |

**Table 473: API Call Metrics for the Admin UI**

| Metric Key | Metric Name | Description |
|---|---|---|
| APICalls\|HTTPRequesterRequestCount | HTTPRequester request count | HTTPRequester request count |
| APICalls\| AvgHTTPRequesterRequestTime | HTTPRequester average request time | HTTPRequester average request time (ms) |

### Suite API Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations API objects.

**Table 474: Suite API Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| UsersCount | Number of users | Number of users |
| ActiveSessionsCount | Active sessions | Active sessions |
| GemfireClientReconnects | Gemfire Client Reconnects | Gemfire Client Reconnects |
| GemfireClientCurrentCalls | Gemfire Client Total Outstanding | Gemfire Client Total Outstanding |
| CurrentHeapSize | Current heap size | Current heap size (MB) . |
| MaxHeapsize | Max heap size | Maximum heap size (MB) . |
| CommittedMemory | Committed memory | Amount of committed memory (MB). |
| CPUUsage | CPU usage | CPU usage (%) . |
| CPUProcessTime | CPU process time | CPU process time (ms) |
| CPUProcessTimeCapacity | CPU process time capacity | CPU process time capacity (ms) |
| Threads | Threads | Number of threads. |

**Table 475: Gemfire Client Call Metrics for the Suite API**

| Metric Key | Metric Name | Description |
|---|---|---|
| GemfireClientCalls\|TotalRequests | Total Requests | Total Requests |
| GemfireClientCalls\|AvgResponseTime | Average Response Time | Average Response Time (ms) |
| GemfireClientCalls\|MinResponseTime | Minimum Response Time | Minimum Response Time (ms) |
| GemfireClientCalls\|MaxResponseTime | Maximum Response Time | Maximum Response Time |
| GemfireClientCalls\|RequestsPerSecond | Requests per Second | Requests per Second |
| GemfireClientCalls\|CurrentRequests | Current Requests | Current Requests |
| GemfireClientCalls\|RequestsCount | Requests Count | Requests Count |
| GemfireClientCalls\|ResponsesCount | Responses Count | Responses Count |

**Table 476: API Call Metrics for the Suite API**

| Metric Key | Metric Name | Description |
|---|---|---|
| APICalls\|TotalRequests | Total Requests | Total Requests |
| APICalls\|AvgResponseTime | Average Response Time (ms) | Average Response Time (ms) |
| APICalls\|MinResponseTime | Minimum Response Time (ms) | Minimum Response Time (ms) |
| APICalls\|MaxResponseTime | Maximum Response Time | Maximum Response Time |
| APICalls\|ServerErrorResponseCount | Server Error Response Count | Server Error Response Count |
| APICalls\|FailedAuthenticationCount | Failed Authentication Count | Failed Authentication Count |
| APICalls\|FailedAuthorizationCount | Failed Authorization Count | Failed Authorization Count |
| APICalls\|RequestsPerSecond | Requests per Second | Requests per Second |
| APICalls\|CurrentRequests | Current Requests | Current Requests |
| APICalls\|ResponsesPerSecond | Responses per Second | Responses per Second |
| APICalls\|RequestsCount | Requests Count | Requests Count |
| APICalls\|ResponsesCount | Responses Count | Responses Count |

### Cluster and Slice Administration Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for VMware Aria OperationsVMware Cloud Foundation Operations Cluster and Slice Administration (CaSA) objects.

**Table 477: Cluster and Slice Administration Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| CurrentHeapSize | Current heap size | Current heap size (MB). |
| MaxHeapsize | Max heap size | Maximum heap size (MB). |
| CommittedMemory | Committed memory | Amount of committed memory (MB). |
| CPUUsage | CPU usage | CPU usage (%) |
| Threads | Threads | Number of threads. |

**Table 478: API Call Metrics for Cluster and Slice Administration**

| Metric Key | Metric Name | Description |
|---|---|---|
| API Calls\|TotalRequests | Total Requests | Total Requests |
| API Calls\|AvgResponseTime | Average Response Time | Average Response Time (ms) |
| API Calls\|MinResponseTime | Minimum Response Time | Minimum Response Time (ms) |
| API Calls\|MaxResponseTime | Maximum Response Time | Maximum Response Time (ms) |
| API Calls\|ServerErrorResponseCount | Server Error Response Count | Server Error Response Count |
| API Calls\|FailedAuthenticationCount | Failed Authentication Count | Failed Authentication Count |
| API Calls\|FailedAuthorizationCount | Minimum Response Time | Minimum Response Time (ms) |

### Watchdog Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects watchdog metrics to ensure that the VMware Aria OperationsVMware Cloud Foundation Operations services are running and responsive.

#### Watchdog Metrics

The watchdog metric provides the total service count.

**Table 479: Watchdog Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| ServiceCount | Service Count | Service Count |

#### Service Metrics

Service metrics provide information about watchdog activity.

**Table 480: Metrics for the VMware Aria OperationsVMware Cloud Foundation Operations Watchdog Service**

| Metric Key | Metric Name | Description |
|---|---|---|
| Service\|Enabled | Enabled | Enabled |
| Service\|Restarts | Restarts | Number of times the process has been unresponsive and been restarted by Watchdog. |
| Service\|Starts | Starts | Number of times the process has been revived by Watchdog. |
| Service\|Stops | Stops | Number of times the process has been stopped by Watchdog. |

### Node Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations node objects.

Metrics can be calculated for node objects. See Calculated Metrics.

**Table 481: Node Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| Component Count | Component count | The number of VMware Aria OperationsVMware Cloud Foundation Operations objects reporting for this node |
| PrimaryResourcesCount | Number of primary objects | Number of primary objects |
| LocalResourcesCount | Number of local objects | Number of local objects |
| PrimaryMetricsCount | Number of primary metrics | Number of primary metrics |
| LocalMetricsCount | Number of local metrics | Number of local metrics |
| PercentDBStorageAvailable | Percent disk available /storage/db | Percent disk available /storage/db |
| PercentLogStorageAvailable | Percent disk available /storage/log | Percent disk available /storage/log |
| FPing stats | Latency\|Average (ms) | |
| FPing stats | Latency\|Maximum (ms) | |
| FPing stats | Latency\|Minimum (ms) | |
| FPing stats | Packet Loss\|Average (%) | Percentage average packet loss between nodes. |
| FPing stats | Packet Loss\|Maximum (%) | Percentage maximum packet loss between nodes. |
| FPing stats | Packet Loss\|Minimum (%) | Percentage minimum packet loss between nodes. |

**Table 482: Memory Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| mem\|actualFree | Actual Free | Actual Free |
| mem\|actualUsed | Actual Used | Actual Used |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| mem\|free | Free | Free ) |
| mem\|used | Used | Used |
| mem\|total | Total | Total |
| mem\|demand_gb | Estimated memory demand | Estimated memory demand |

**Table 483: Swap Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| swap\|total | Total | Total |
| swap\|free | Free | Free |
| swap\|used | Used | Used |
| swap\|pageIn | Page in | Page in |
| swap\|pageOut | Page out | Page out |

**Table 484: Resource Limit Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| resourceLimit\|numProcesses | Number of processes | Number of processes |
| resourceLimit\|openFiles | Number of open files | Number of open files |
| resourceLimit\|openFilesMax | Number of open files maximum limit | Number of open files maximum limit |
| resourceLimit\|numProcessesMax | Number of processes maximum limit | Number of processes maximum limit |

**Table 485: Network Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| net\|allInboundTotal | All inbound connections | All inbound total |
| net\|allOutboundTotal | All outbound connections | All outbound total |
| net\|tcpBound | TCP bound | TCP bound |
| net\|tcpClose | TCP state CLOSE | Number of connections in TCP CLOSE |
| net\|tcpCloseWait | TCP state CLOSE WAIT | Number of connections in TCP state CLOSE WAIT |
| net\|tcpClosing | TCP state CLOSING | Number of connections in TCP state CLOSING |
| net\|tcpEstablished | TCP state ESTABLISHED | Number of connections in TCP state ESTABLISHED |
| net\|tcpIdle | TCP state IDLE | Number of connections in TCP state IDLE |
| net\|tcpInboundTotal | TCP inbound connections | TCP inbound connections |
| net\|tcpOutboundTotal | TCP outbound connections | TCP outbound connections |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| net\|tcpLastAck | TCP state LAST ACK | Number of connections in TCP state LAST ACK |
| net\|tcpListen | TCP state LISTEN | Number of connections in TCP state LISTEN |
| net\|tcpSynRecv | TCP state SYN RCVD | Number of connections in TCP state SYN RCVD |
| net\|tcpSynSent | TCP state SYN_SENT | Number of connections in TCP state SYN_SENT |
| net\|tcpTimeWait | TCP state TIME WAIT | Number of connections in TCP state TIME WAIT |

**Table 486: Network Interface Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| net\|iface\|speed | Speed | Speed (bits/sec) |
| net\|iface\|rxPackets | Receive packets | Number of received packets |
| net\|iface\|rxBytes | Receive bytes | Number of received bytes |
| net\|iface\|rxDropped | Receive packet drops | Number of received packets dropped |
| net\|iface\|rxFrame | Receive packets frame | Number of receive packets frame |
| net\|iface\|rxOverruns | Receive packets overruns | Number of receive packets overrun |
| net\|iface\|txPackets | Transmit packets | Number of transmit packets |
| net\|iface\|txBytes | Transmit bytes | Number of transmit bytes |
| net\|iface\|txDropped | Transmit packet drops | Number of transmit packets dropped |
| net\|iface\|txCarrier | Transmit carrier | Transmit carrier |
| net\|iface\|txCollisions | Transmit packet collisions | Number of transmit collisions |
| net\|iface\|txErrors | Transmit packet errors | Number of transmit errors |
| net\|iface\|txOverruns | Transmit packet overruns | Number of transmit overruns |

**Table 487: Disk Filesystem Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| disk\|fileSystem\|total | Total | Total |
| disk\|fileSystem\|available | Available | Available |
| disk\|fileSystem\|used | Used | Used |
| disk\|fileSystem\|files | Total file nodes | Total file nodes |
| disk\|fileSystem\|filesFree | Total free file nodes | Total free file nodes |
| disk\|fileSystem\|queue | Disk queue | Disk queue |
| disk\|fileSystem\|readBytes | Read bytes | Number of bytes read |
| disk\|fileSystem\|writeBytes | Write bytes | Number of bytes written |
| disk\|fileSystem\|reads | Reads | Number of reads |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| disk\|fileSystem\|writes | Writes | Number of writes |

**Table 488: Disk Installation Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| disk\|installation\|used | Used | Used |
| disk\|installation\|total | Total | Total |
| disk\|installation\|available | Available | Available |

**Table 489: Disk Database Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| disk\|db\|used | Used | Used |
| disk\|db\|total | Total | Total |
| disk\|db\|available | Available | Available |

**Table 490: Disk Log Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| disk\|log\|used | Used | Used |
| disk\|log\|total | Total | Total |
| disk\|log\|available | Available | Available |

**Table 491: CPU Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| cpu\|combined | Combined load | Combined load (User + Sys + Nice + Wait) |
| cpu\|idle | Idle | Idle time fraction of total available cpu (cpu load) |
| cpu\|irq | Irq | Interrupt time fraction of total available cpu (cpu load) |
| cpu\|nice | Nice | Nice time fraction of total available cpu (cpu load) |
| cpu\|softIrq | Soft Irq | Soft interrupt time fraction of total available cpu (cpu load) |
| cpu\|stolen | Stolen | Stolen time fraction of total available cpu (cpu load) |
| cpu\|sys | Sys | Sys time fraction of total available cpu (cpu load) |
| cpu\|user | User (cpu load) | User time fraction of total available cpu (cpu load) |
| cpu\|wait | Wait (cpu load) | Wait time fraction of total available cpu (cpu load) |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| cpu\|total | Total available for a cpu | Total available for a cpu |
| cpu\|allCpuCombined | Total combined load for all cpus | Total combined load for all cpus (cpu load) |
| cpu\|allCpuTotal_ghz | Available | Available |
| cpu\|allCpuCombined_ghz | Used | Used |
| cpu\|allCpuCombined_percent | CPU usage | CPU usage (%) |

**Table 492: Device Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| device\|iops | Reads/Writes per second | Average number of read/write commands issued per second during the collection interval. |
| device\|await | Average transaction time | Average transaction time (milliseconds). |
| device\|iops_readMaxObserved | Maximum observed reads per second | Maximum observed reads per second. |
| device\|iops_writeMaxObserved | Maximum observed writes per second | Maximum observed writes per second. |

**Table 493: Service Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| service\|proc\|fdUsage | Total number of open file descriptors | Total number of open file descriptors. |

**Table 494: NTP Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| ntp\|serverCount | Configured server count | Configured server count |
| ntp\|unreachableCount | Unreachable server count | Unreachable server count |
| ntp\|unreachable | Unreachable | Is the NTP server unreachable. Value of 0 is reachable, 1 means the server was not reached or did not respond. |

**Table 495: Heap Metrics for the Node**

| Metric Key | Metric Name | Description |
|---|---|---|
| heap\|CurrentHeapSize | Current heap size | Current heap size |
| heap\|MaxHeapSize | Max heap size | Max heap size |
| heap\|CommittedMemory | Committed Memory | Committed Memory |

### Cluster Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the VMware Aria OperationsVMware Cloud Foundation Operations cluster objects including dynamic threshold calculation metrics and capacity computation metrics.

Metrics can be calculated for cluster objects. See Calculated Metrics.

#### Cluster Metrics

Cluster metrics provide host, resource, and metric counts on the cluster.

**Table 496: Cluster Metrics**

| Metric Key | Metric Name | Description |
|---|---|---|
| HostCount | Number of Nodes in Cluster | Number of Nodes in Cluster |
| PrimaryResourcesCount | Number of primary resources | Number of primary resources |
| LocalResourcesCount | Number of local resources | Number of local resources |
| PrimaryMetricsCount | Number of primary metrics | Number of primary metrics |
| ReceivedResourceCount | Number of received resources | Number of received resources |
| ReceivedMetricCount | Number of received metrics | Number of received metrics |
| Count (VM) | Usage Count | This metric displays how many units of the license capacity is currently used. |
| License Metrics for Cluster Object | | |
| Used (VM) | Usage (%) | This metric displays the percentage of the total license capacity currently used.<br>**NOTE**<br>An alert is generated if the license threshold is in any of the following states:<br>• >= 80% - Warning<br>• >= 90% - Immediate<br>• =95% - Catastrophic |
| Days Remaining (day) | Days Remaining | This metric displays the days remaining before the license expires. |
| Capacity | Capacity | Displays the maximum number of units (of the given capacity type) that can be licensed by this license key. |
| Type | Type | Displays the license type for the cluster object. |
| Expiry | Expiration Date | Displays the date when the license expires. |
| Capacity | Capacity | Displays the maximum number of units (of the given capacity type) that can be licensed by this license key. |
| Disbalance factor (%) | Disbalance factor (%) | Identifies the state of disbalance in a VMware Aria OperationsVMware |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| | | Cloud Foundation Operations cluster and calculates the disbalance factor (when one or more nodes have a higher density among all the nodes in the cluster). Based on the disbalance percentage of the node, the following alerts are triggered:<br>• Warning: If the disbalance factor metric is equal to or greater than 5%.<br>• Immediate: If the disbalance factor metric is equal to or greater than 7%.<br>• Critical: If the disbalance factor metric is equal to or greater than 10%. |

**DT Metrics**

DT metrics are dynamic threshold metrics for the cluster. Non-zero values appear only if metric collection occurs while the dynamic threshold calculations are running.

**Table 497: DT Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| dt\|isRunning | Running | Running |
| dt\|dtRunTime | Running duration | Running duration (ms) |
| dt\|StartTime | Running start time | Running start time |
| dt\|percentage | Percent | Percent (%) |
| dt\|executorCount | Executor Node Count | Executor Node Count |
| dt\|resourceCount | Resource Count | Resource Count |
| dt\|fsdbReadTime | FSDB Read Time | FSDB Read Time (ms) |
| dt\|dtObjectSaveTime | DT Object Save Time | DT Object Save Time (ms) |
| dt\|dtHistorySaveTime | DT History Save Time | DT History Save Time (ms) |
| dt\|executor\|resourceCount | Resource Count | Resource Count |

**Capacity Computation (CC) Metrics**

CC metrics are capacity computation metrics for the cluster. Non-zero values appear only if metric collection occurs while the capacity computation calculations are running.

**Table 498: CC Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| cc\|isRunning | Running | Running |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| cc\|runTime | Total Run Time | Total Run Time |
| cc\|startTime | Start time | Start time |
| cc\|finishTime | Finish Time | Finish Time |
| cc\|totalResourcesToProcess | Total Objects Count | Total Objects Count |
| cc\|progress | Progress | Progress |
| cc\|phase1TimeTaken | Phase 1 Computation Time | Phase 1 Computation Time |
| cc\|phase2TimeTaken | Phase 2 Computation Time | Phase 2 Computation Time |

**Gemfire Cluster Metrics**

Gemfire metrics provide information about the Gemfire cluster.

**Table 499: Gemfire cluster Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| GemfireCluster\|System\|AvgReads | Average reads per second | The average number of reads per second for all members |
| GemfireCluster\|System\|AvgWrites | Average writes per second | The average number of writes per second for all members |
| GemfireCluster\|System\|DiskReadsRate | Disk reads rate | The average number of disk reads per second across all distributed members |
| GemfireCluster\|System\|DiskWritesRate | Disk writes rate | The average number of disk writes per second across all distributed members |
| GemfireCluster\|System\|GarbageCollectionCount | Total garbage collection count | The total garbage collection count for all members |
| GemfireCluster\|System\|GarbageCollectionCountDelta | New garbage collection count | The new garbage collection count for all members |
| GemfireCluster\|System\|JVMPauses | JVM pause count | The number of detected JVM pauses |
| GemfireCluster\|System\|JVMPausesDelta | New JVM pause count | The number of new detected JVM pauses |
| GemfireCluster\|System\|DiskFlushAvgLatency | Disk flush average latency | Disk flush average latency (msec) |
| GemfireCluster\|System\|NumRunningFunctions | Number of running functions | The number of map-reduce jobs currently running on all members in the distributed system |
| GemfireCluster\|System\|NumClients | Number of clients | The number of connected clients |
| GemfireCluster\|System\|TotalHitCount | Total hit count | Total number of cache hits for all regions |
| GemfireCluster\|System\|TotalHitCountDelta | New hit count | Number of new cache hits for all regions |
| GemfireCluster\|System\|TotalMissCount | Total miss count | The total number of cache misses for all regions |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| GemfireCluster\|System\|TotalMissCountDelta | New miss count | Number of new cache misses for all regions |
| GemfireCluster\|System\|Member\|FreeSwapSpace | Swap space free | Swap space free (MB) |
| GemfireCluster\|System\|Member\|TotalSwapSpace | Swap space total | Swap space total (MB) |
| GemfireCluster\|System\|Member\|CommittedVirtualMemorySize | Committed virtual memory size | Committed virtual memory size (MB) |
| GemfireCluster\|System\|Member\|SystemLoadAverage | System load average | System load average |
| GemfireCluster\|System\|Member\|FreePhysicalMemory | Free physical memory | Free physical memory (MB) |
| GemfireCluster\|System\|Member\|TotalPhysicalMemory | Total physical memory | Total physical memory (MB) |
| GemfireCluster\|System\|Member\|CacheListenerCallsAvgLatency | Average cache listener calls latency | Average cache listener calls latency (msec) |
| GemfireCluster\|System\|Member\|CacheWriterCallsAvgLatency | Average cache writer calls latency | Average cache writer calls latency (msec) |
| GemfireCluster\|System\|Member\|DeserializationAvgLatency | Average deserialization latency | Average deserialization latency (msec) |
| GemfireCluster\|System\|Member\|FunctionExecutionRate | Function executions per second | Function executions per second |
| GemfireCluster\|System\|Member\|JVMPauses | Number of JVM pauses | Number of JVM pauses |
| GemfireCluster\|System\|Member\|NumRunningFunctions | Number of running functions | Number of running functions |
| GemfireCluster\|System\|Member\|PutsRate | Puts per second | Puts per second |
| GemfireCluster\|System\|Member\|GetsRate | Gets per second | Gets per second |
| GemfireCluster\|System\|Member\|GetsAvgLatency | Average gets latency | Average gets latency (msec) |
| GemfireCluster\|System\|Member\|PutsAvgLatency | Average puts latency | Average puts latency (msec) |
| GemfireCluster\|System\|Member\|SerializationAvgLatency | Average serialization latency | Average serialization latency (msec) |
| GemfireCluster\|System\|Member\|Disk\|DiskFlushAvgLatency | Flush average latency | Flush average latency (msec) |
| GemfireCluster\|System\|Member\|Disk\|DiskReadsRate | Average reads per second | Average reads per second |
| GemfireCluster\|System\|Member\|Disk\|DiskWritesRate | Average writes per second | Average writes per second |
| GemfireCluster\|System\|Member\|Network\|BytesReceivedRate | Average received bytes per second | Average received bytes per second |
| GemfireCluster\|System\|Member\|Network\|BytesSentRate | Average sent bytes per second | Average sent bytes per second |
| GemfireCluster\|System\|Member\|JVM\|GCTimeMillis | Garbage Collection time | Total amount of time spent on garbage collection |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| GemfireCluster\|System\|Member\|JVM\|GCTimeMillisDelta | New Garbage Collection time | New amount of time spent on garbage collection |
| GemfireCluster\|System\|Member\|JVM\|TotalThreads | Total threads | Total threads |
| GemfireCluster\|System\|Member\|JVM\|CommitedMemory | Committed Memory | Committed Memory (MB) |
| GemfireCluster\|System\|Member\|JVM\|MaxMemory | Max Memory | Max Memory (MB) |
| GemfireCluster\|System\|Member\|JVM\|UsedMemory | Used Memory | Used Memory (MB) |
| GemfireCluster\|Region\|SystemRegionEntryCount | Entry Count | Entry Count |
| GemfireCluster\|Region\|DestroyRate | Destroys per second | Destroys per second |
| GemfireCluster\|Region\|CreatesRate | Creates per second | Creates per second |
| GemfireCluster\|Region\|GetsRate | Gets per second | Gets per second |
| GemfireCluster\|Region\|BucketCount | Bucket count | Bucket count |
| GemfireCluster\|Region\|AvgBucketSize | Average number of entries per bucket | Average number of entries per bucket |
| GemfireCluster\|Region\|Member\|ActualRedundancy | Actual redundancy | Actual redundancy |
| GemfireCluster\|Region\|Member\|BucketCount | Bucket count | Bucket count |
| GemfireCluster\|Region\|Member\|AvgBucketSize | Average number of entries per bucket | Average number of entries per bucket |
| GemfireCluster\|Region\|Member\|CreatesRate | Creates per second | Creates per second |
| GemfireCluster\|Region\|Member\|GetsRate | Gets per second | Gets per second |
| GemfireCluster\|Region\|Member\|DestroyRate | Destroys per second | Destroys per second |
| GemfireCluster\|Region\|Member\|MissCount | Number of misses count | Number of cache misses |
| GemfireCluster\|Region\|Member\|MissCountDelta | Number of new cache misses | Number of new cache misses |
| GemfireCluster\|Region\|Member\|HitCount | Number of hits count | Number of cache hits |
| GemfireCluster\|Region\|Member\|HitCountDelta | Number of new cache hits | Number of new cache hits |

## Threshold Checking Metrics

Threshold checking metrics check the processed and computed metrics for the cluster.

**Table 500: Threshold Checking Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| ThresholdChecking\|ProcessedMetricCount | Number of processed metrics | Number of processed metrics |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| ThresholdChecking\|ProcessedMetricRate | Received metric processing rate (per second) | Received metric processing rate (per second) |
| ThresholdChecking\|ComputedMetricCount | Number of computed metrics | Number of computed metrics |
| ThresholdChecking\|ComputedMetricRate | Computed metric processing rate (per second) | Computed metric processing rate (per second) |

**Memory Metrics**

Memory metrics provide memory CPU use information for the cluster.

**Table 501: Memory Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| Memory\|AvgFreePhysicalMemory | Average free physical memory | Average free physical memory (GB) |
| Memory\|TotalFreePhysicalMemory | Free physical memory | Free physical memory (GB) |
| Memory\|TotalMemory | Total Available Memory | Total Available Memory (GB) |
| Memory\|TotalUsedMemory | Actual Used Memory | Actual Used Memory (GB) |
| Memory\|TotalDemandMemory | Memory Demand | Memory Demand (GB) |

**Elastic Memory Metrics**

Elastic memory metrics provide reclaimable memory CPU use information for the cluster.

**Table 502: Memory Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| ElasticMemory\|TotalMemory | Total Available Memory | Total Available Memory (GB) |
| ElasticMemory\|TotalUsedMemory | Actual Used Memory | Actual Used Memory (GB) |
| ElasticMemory\|TotalDemandMemory | Memory Demand | Memory Demand (GB) |

**CPU Metrics**

CPU metrics provide CPU information for the cluster.

**Table 503: CPU Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| cpu\|TotalCombinedUsage | CPU Load | CPU Load |
| cpu\|TotalAvailable | CPU Available | CPU Available |
| cpu\|TotalAvailable_ghz | Available | Available (GHz) |
| cpu\|TotalUsage_ghz | Used | Used (GHz) |
| cpu\|TotalUsage | CPU usage | CPU usage (%) |

### Disk Metrics

Disk metrics provide available disk information for the cluster.

**Table 504: Disk Metrics for the Cluster**

| Metric Key | Metric Name | Description |
|---|---|---|
| Disk\|DatabaseStorage\|AvgAvailable | Average node disk available | Average node disk available |
| Disk\|DatabaseStorage\|MinAvailable | Minimum node disk available | Minimum node disk available |
| Disk\|DatabaseStorage\|MaxAvailable | Maximum node disk available | Maximum node disk available |
| Disk\|DatabaseStorage\|TotalAvailable | Available | Available |
| Disk\|DatabaseStorage\|Total | Total | Total |
| Disk\|DatabaseStorage\|TotalUsed | Used | Used |
| Disk\|LogStorage\|AvgAvailable | Average node disk available | Average node disk available |
| Disk\|LogStorage\|MinAvailable | Minimum node disk available | Minimum node disk available |
| Disk\|LogStorage\|MaxAvailable | Maximum node disk available | Maximum node disk available |
| Disk\|LogStorage\|TotalAvailable | Available | Available |
| Disk\|LogStorage\|Total | Total | Total |
| Disk\|LogStorage\|TotalUsed | Used | Used |

## Persistence Metrics

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for various persistence resources or service groups.

### Activity Metrics

Activity metrics relate to the activity framework.

**Table 505: Activity Metrics for Persistence**

| Metric Key | Metric Name | Description |
|---|---|---|
| Activity\|RunningCount | Number Running | Number Running |
| Activity\|ExecutedCount | Number Executed | Number Executed |
| Activity\|SucceededCount | Number Succeeded | Number Succeeded |
| Activity\|FailedCount | Number Failed | Number Failed |

### Controller XDB Metrics

Controller metrics relate to the primary database.

**Table 506: Controller XDB Metrics for Persistence**

| Metric Key | Metric Name | Description |
|---|---|---|
| ControllerXDB\|Size | Size | Size (Bytes) |
| ControllerXDB\|TempDBSize | Temporary DB Size | Temporary DB Size (Bytes) |
| ControllerXDB\|TotalObjectCount | Total Object Count | Total Object Count |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| ControllerXDB\|AvgQueryDuration | Average Query Duration | Average Query Duration (ms) |
| ControllerXDB\|MinQueryDuration | Minimum Query Duration | Minimum Query Duration (ms) |
| ControllerXDB\|MaxQueryDuration | Maximum Query Duration | Maximum Query Duration (ms) |
| ControllerXDB\|TotalTransactionCount | Total Transaction Count | Total Transaction Count |
| ControllerXDB\|LockOperationErrorCount | Lock Operation Error Count | Lock Operation Error Count |
| ControllerXDB\|DBCorruptionErrorCount | DB Corruption Error Count | DB Corruption Error Count |
| ControllerXDB\| DBMaxSessionExceededCount | DB Maximum Sessions Exceeded Count | DB Maximum Sessions Exceeded Count |
| ControllerXDB\|NumberWaitingForSession | Number of operations waiting for a session | Number of operations waiting for a session from the session pool |
| ControllerXDB\|AvgWaitForSessionDuration | Average acquisition time from session pool | Average acquisition time from session pool |
| ControllerXDB\|MinWaitForSessionDuration | Minimum acquisition time from session pool | Minimum acquisition time from session pool |
| ControllerXDB\|MaxWaitForSessionDuration | Maximum acquisition time from session pool | Maximum acquisition time from session pool |
| ControllerXDB\|TotalGetSessionCount | Total requests for a session from the session pool | Total requests for a session from the session pool |
| ControllerXDB\|MaxActiveSessionCount | Maximum Concurrent Session Count | Maximum concurrent session count during the past collection interval. |

**Alarm SQL Metrics**

Alarm metrics relate to the persistence of alerts and symptoms.

**Table 507: Alarm XDB Metrics for Persistence**

| Metric Key | Metric Name | Description |
|---|---|---|
| AlarmSQL\|Size | Size (Bytes) | Size (Bytes) |
| AlarmSQL\|AvgQueryDuration | Average Query Duration (ms) | Average Query Duration (ms) |
| AlarmSQL\|MinQueryDuration | Minimum Query Duration (ms) | Minimum Query Duration (ms) |
| AlarmSQL\|MaxQueryDuration | Maximum Query Duration (ms) | Maximum Query Duration (ms) |
| AlarmSQL\|TotalTransactionCount | Total Transaction Count | Total Transaction Count |
| AlarmSQL\|TotalAlarms | Alarm Total Object Count | Alarm Total Object Count |
| AlarmSQL\|TotalAlerts | Alert Total Object Count | Alert Total Object Count |
| AlarmSQL\|AlertTableSize | Alert Table Size | Alert Table Size |
| AlarmSQL\|AlarmTableSize | Alarm Table Size | Alarm Table Size |

**Key Value Store Database (KVDB)**

KVDB metrics relate to the persistence of storing key-value data.

| Metric Key | Metric Name | Description |
|---|---|---|
| KVDB\|AvgQueryDuration | Average Query Duration | Average Query Duration |
| KVDB\|MinQueryDuration | Minimum Query Duration | Minimum Query Duration |
| KVDB\|MaxQueryDuration | Maximum Query Duration | Maximum Query Duration |
| KVDB\|TotalTransactionCount | Total Transaction Count | Total Transaction Count |

## Historical Inventory Service XDB Metrics

Historical inventory service metrics relate to the persistence of configuration properties and their changes.

**Table 508: Historical XDB Metrics for Persistence**

| Metric Key | Metric Name | Description |
|---|---|---|
| HisXDB\|FunctionCalls\|Count HisXDB\|FunctionCalls | Number of Function calls | Number of Function calls |
| HisXDB\|FunctionCalls\|AvgDuration | Average execution time | Average execution time |
| HisXDB\|FunctionCalls\|MaxDuration | Max execution time | Max execution time |
| HisXDB\|Size | Size | Size (Bytes) |
| HisXDB\|TempDBSize | Temporary DB Size | Temporary DB Size (Bytes) |
| HisXDB\|TotalObjectCount | Total Object Count | Total Object Count |
| HisXDB\|AvgQueryDuration | Average Query Duration | Average Query Duration (ms) |
| HisXDB\|MinQueryDuration | Minimum Query Duration | Minimum Query Duration (ms) |
| HisXDB\|MaxQueryDuration | Maximum Query Duration | Maximum Query Duration (ms) |
| HisXDB\|TotalTransactionCount | Total Transaction Count | Total Transaction Count |
| HisXDB\|LockOperationErrorCount | Lock Operation Error Count | Lock Operation Error Count |
| HisXDB\|DBCorruptionErrorCount | DB Corruption Error Count | DB Corruption Error Count |
| HisXDB\|DBMaxSessionExceededCount | DB Maximum Sessions Exceeded Count | DB Maximum Sessions Exceeded Count |
| HisXDB\|NumberWaitingForSession | Number of operations waiting for a session | Number of operations waiting for a session from the session pool |
| HisXDB\|AvgWaitForSessionDuration | Average acquisition time from session pool | Average acquisition time from session pool |
| HisXDB\|MinWaitForSessionDuration | Minimum acquisition time from session pool | Minimum acquisition time from session pool |
| HisXDB\|MaxWaitForSessionDuration | Maximum acquisition time from session pool | Maximum acquisition time from session pool |
| HisXDB\|TotalGetSessionCount | Total requests for a session from the session pool | Total requests for a session from the session pool |
| HisXDB\|HisActivitySubmissionCount | HIS activity submission count | Number of Historical Inventory Service activities submitted |
| HisXDB\|HisActivityCompletionCount | HIS activity completion count | Number of Historical Inventory Service activities completed |
| HisXDB\|HisActivityCompletionDelayAvg | HIS activity average completion delay | The average amount of time from activity submission to completion |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| HisXDB\|HisActivityCompletionDelayMax | HIS activity maximum completion delay | The maximum amount of time from activity submission to completion |
| HisXDB\|HisActivityAbortedCount | HIS activity abort count | Number of Historical Inventory Service activities stopped |

### VMware Aria Automation Metrics

VMware Aria Automation collects metrics for objects such as, cloud zone, project, deployment, blueprint, cloud account, user, and cloud automation services world Instance.

### Blueprint Metrics

VMware Aria Automation collects metrics for objects such as blueprint object.

**Table 509: Blueprint Metrics**

| Property Name | Metrics |
|---|---|
| Summary | VMCount |

### Project Metrics

VMware Aria Automation collects metrics for objects such as project object.

**Table 510: Project Metrics**

| Property Name | Metrics |
|---|---|
| Summary | VMCount |
| Summary | TotalDeployments |
| Summary | TotalCloudZones |
| Summary | TotalBlueprints |
| Summary | Metering\|Additional price |
| Summary | Metering\|CPU Price |
| Summary | Metering\|Memory price |
| Summary | Metering\|Storage Price |
| Summary | Metering\|Total price |

### Deployment Metrics

VMware Aria Automation collects the metrics for the deployment object.

**Table 511: Deployment Metrics**

| Property Name | Metrics |
|---|---|
| Summary | Metering\|Additional price |
| Summary | Metering\|CPU Price |
| Summary | Metering\|Memory price |
| Summary | Metering\|Storage Price |
| Summary | Metering\|Total price |
| Summary | Metering\|Partial price |

### Organization Metrics

VMware Aria Automation collects the metrics for the organization object.

**Table 512: Organization Metrics**

| Property Name | Metrics |
|---|---|
| Summary | TotalBlueprints |
| Summary | TotalProjects |
| Summary | VMCount |
| Summary | TotalDeployments |
| Summary | TotalCloudZones |

### VMware Aria Automation Adapter Metrics

VMware Aria Automation collects the metrics for the adapter object.

**Table 513: VMware Aria Automation Adapter Metrics**

| Property Name | Metrics |
|---|---|
| Summary | TotalCloudZones |
| Summary | VMCount |
| Summary | TotalDeployments |
| Summary | TotalBlueprints |
| Summary | TotalProjects |

### Cloud Automation Services World Metrics

VMware Aria Automation collects the metrics for the Cloud Automation Services world object.

**Table 514: Cloud Automation Services World Metrics**

| Property Name | Metrics |
|---|---|
| Summary | TotalDeployments |

*Table continued on next page*

*Continued from previous page*

| Property Name | Metrics |
|---|---|
| Summary | VMCount |
| Summary | TotalCloudZones |
| Summary | TotalProjects |
| Summary | TotalBlueprints |

### Cloud Automation Services Entity Status Metrics

VMware Aria Automation collects the metrics for the Cloud Automation Services (CAS) entity status object.

**Table 515: Cloud Automation Services Entity Status Metrics**

| Property Name | Metrics |
|---|---|
| Summary | TotalClusters |

### Cloud Zone Metrics

VMware Aria Automation collects memory, storage, and vCPU limit related metrics at project level on per cloud zone basis.

#### Cloud Zone Limits Metrics

You can view the memory, storage, and vCPU limit related metrics defined in VMware Aria Automation in VMware Aria OperationsVMware Cloud Foundation Operations. VMware Aria OperationsVMware Cloud Foundation Operations alerts when the utilisation exceeds configured limits in any of the individual cloud zones.

> **NOTE**
> The resource limits are configured in VMware Aria Automation. For vCenter objects, limits for all resources - vCPU, memory, and storage - can be configured. However, for public clouds like AWS, Azure, and GCP only vCPU and memory limits can be configured.

**Table 516: Cloud Zone Limits Metrics**

| Property Name | Metrics | Description |
|---|---|---|
| Cloud Zone Limits | Memory allocated (KB) | This metric displays the total utilized cloud zone memory. |
| | Memory Limit (KB) | This metric displays the limit set on the cloud zone memory usage. |
| | Memory Utilized (%) | This metric displays the percentage of memory used out of the memory limit set. Formula ((Memory allocated (KB)/ Memory Limit (KB)) |
| | Storage allocated (GB) | This metric displays the total storage utilized by the cloud zone. |
| | Storage Limit (GB) | This metrics displays the limit set on the cloud zone storage usage. |

*Table continued on next page*

*Continued from previous page*

| Property Name | Metrics | Description |
|---|---|---|
| | Storage Utilized (%) | This metric displays the percentage of storage used out of the storage limit set. Formula (Storage allocated (KB)/ Storage Limit (KB)) |
| | vCPU allocated | This metric displays the total vCPU utilized by the cloud zone. |
| | vCPU Limit | This metrics displays the limit set on the cloud zone vCPU usage. |
| | vCPU Utilized (%) | This metric displays the percentage of vCPU used out of the vCPU limit set. Formula (vCPU allocated (KB)/ vCPU Limit (KB)) |

**Summary Metrics for Cloud Zone Centers**

Summary metrics provide information about the VM migration as part of workload planning.

| Metric Name | Description |
|---|---|
| WLP | Displays the VM migration trend as part of workload optimization. These metrics are deactivated by default. You must activate them from policies.<br>• Fail Count: Number of failed VM move attempts in the last daily cycle.<br>• Number of runs: The total number of times the WLP was run during the last daily cycle.<br>• Success Count: The number of successful VM moves during the last daily cycle. |

### Metrics for vSAN

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for vSAN objects.

From the left menu, click **Inventory** > **Inventory Panel (Detailed View)** > **Integrations** > **All Objects** > **vSAN Adapter**, and then select one of the vSAN adapter objects listed and click the **Metrics** tab.

#### Disk I/O and Disk Space Metrics for vSAN Disk Groups

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN disk groups.

Disk I/O metrics for the vSAN disk groups include:

- Disk I/O|Reads Per Second (IOPS)
- Disk I/O|Writes Per Second (IOPS
- Disk I/O|Max Observed Reads Per Second (IOPS)
- Disk I/O|Max Observed Writes Per Second (IOPS)
- Disk I/O|Throughput Read (bps)
- Disk I/O|Throughput Write (bps)

- Disk I/O|Average Read Latency (ms)
- Disk I/O|Average Write Latency (ms)
- Disk I/O|Total Bus Resets
- Disk I/O|Total Commands Aborted per second

The following Disk I/O metrics are disabled by default:

- Disk I/O|Read Count
- Disk I/O|Write Count
- Disk I/O|Average Device Latency
- Disk I/O|Average Device Read Latency
- Disk I/O|Average Device Write Latency
- Disk I/O|Total Number of Errors

Disk space metrics for vSAN disk groups include:

- Disk Space|Capacity (bytes)
- Disk Space|Used (bytes)
- Disk Space|Usage (%)

### Read Cache Metrics for vSAN Disk Groups

The VMware Aria OperationsVMware Cloud Foundation Operations collects metrics and performs capacity trend analysis on a hybrid vSAN read cache. Read Cache metrics are not collected for a vSAN all-flash configuration.

Read cache metrics for the vSAN disk group include:

- Read Cache|Hit Rate (%)
- Read Cache|Miss Rate Ratio
- Read Cache|Reads Per Second (IOPS)
- Read Cache|Read Latency (ms)
- Read Cache|Writes Per Second (IOPS)
- Read Cache|Write Latency (ms)

The following read cache metrics are disabled by default:

- Read Cache|Read I/O Count
- Read Cache|Write I/O Count

### Write Buffer Metrics for vSAN Disk Groups

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the write buffer capacity of your vSAN disk groups.

A reasonably balanced system consumes a significant amount of write buffer. Before placing additional workload on the vSAN, check the write buffer metrics for the vSAN disk group.

- Write Buffer|Capacity (bytes)
- Write Buffer|Free (%)
- Write Buffer|Usage (%)
- Write Buffer|Used (byte)
- Write Buffer|Reads Per Second (IOPS)
- Write Buffer|Read Latency (ms)
- Write Buffer|Writes Per Second (IOPS)
- Write Buffer|Write Latency (ms)

The following write buffer metrics are disabled by default:

- Write Buffer|Read I/O Count
- Write Buffer|Write I/O Count

### Congestion Metrics for vSAN Disk Groups

The VMware Aria OperationsVMware Cloud Foundation Operations collects congestion metrics for the vSAN disk group.

- Congestion| Memory Congestion - Favorite
- Congestion| SSD Congestion - Favorite
- Congestion| IOPS Congestion - Favorite
- Congestion| Slab Congestion
- Congestion| Log Congestion
- Congestion| Comp Congestion

### Cache De-stage Metrics for vSAN Disk Groups

The VMware Aria OperationsVMware Cloud Foundation Operations collects cache de-stage metrics for the vSAN disk groups.

Cache de-stage metrics include:

- Bytes De-stage from SSD
- Zero-bytes De-stage

### Resync Traffic Metrics for vSAN Disk Groups

The VMware Aria OperationsVMware Cloud Foundation Operations collects resync traffic metrics for the vSAN disk groups.

Resync traffic metrics include:

- Read IOPS for Resync Traffic
- Write IOPS for Resync Traffic
- Read Throughput for Resync Traffic
- Write Throughput for Resync Traffic
- Read Latency for Resync Traffic
- Write Latency for Resync Traffic

### Metrics for vSAN Cluster

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN cluster.

VMware Aria OperationsVMware Cloud Foundation Operations enhances the capacity calculation for vSAN, using the new slack space provided by the new vSAN API. Cost calculation is still done using the old way which reserves 30% memory for Slack Overhead.

Metrics for vSAN cluster include:

| Component | Metrics |
|---|---|
| Component Limit | • vSAN\|Component Limit\|Component Limit Used (%)<br>• vSAN\|Component Limit\|Total Component Limit<br>• vSAN\|Component Limit\|Used Component Limit |
| Disk Space | • vSAN\|Disk Space\|Disk Space Used (%)<br>• vSAN\|Disk Space\|Total Disk Space (GB)<br>• vSAN\|Disk Space\|Used Disk Space (GB)<br>• vSAN\|Disk Space\|Usable Capacity (GB)<br>• vSAN\|Disk Space\|Oversubscription Capacity (GB)<br>• vSAN\|Disk Space\|Oversubscription Ratio<br>**NOTE**<br>The oversubscription metrics are applicable only for the vSAN OSA clusters. |
| Health | vSAN\|Health\|Score |
| Read Cache | • vSAN\|Read Cache\|Read Cache Reserved (%)<br>• vSAN\|Read Cache\|Reserved Read Cache Size (GB)<br>• vSAN\|Read Cache\|Total Read Cache Size (GB) |
| Performance | • vSAN\|Read Cache\|Reads Per Second (IOPS)<br>• vSAN\|Read Cache\|Read Throughput (KBps)<br>• vSAN\|Read Cache\|Average Read Latency (ms)<br>• vSAN\|Read Cache\|Writes Per Second (IOPS)<br>• vSAN\|Read Cache\|Write Throughput (KBps)<br>• vSAN\|Read Cache\|Average Write Latency (ms)<br>• vSAN\|Read Cache\|Congestion<br>• vSAN\|Read Cache\|Outstanding I/O<br>• vSAN\|Read Cache\|Total IOPS<br>• vSAN\|Read Cache\|Total Latency (ms)<br>• vSAN\|Read Cache\|Total Throughput (KBps) |
| Deduplication And Compression Overview | • vSAN\|Deduplication And Compression Overview\|Used Before<br>• vSAN\|Deduplication And Compression Overview\|Used After<br>• vSAN\|Deduplication And Compression Overview\|Savings<br>• vSAN\|Deduplication And Compression Overview\|Ratio |
| Summary | • Summary\|Number of Storage Pools<br>• Summary\|Number of ESA Disks<br>• Summary\|Number of Cache Disks<br>• Summary\|Total Number of Capacity Disks<br>• Summary\|CPU Workload<br>• Summary\|Memory Workload<br>• Summary\|Total Number of Disk Groups<br>• Summary\|Total Active Alerts Count<br>• Summary\|Total Number of VMs<br>• Summary\|Total Number of Hosts<br>• Summary\|vSAN Cluster Capacity Remaining (%)<br>• Summary\|vSAN Cluster Storage Time Remaining |

*Table continued on next page*

*Continued from previous page*

| Component | Metrics |
|---|---|
| | • Summary\|vSAN Capacity Disk Used<br>• Summary \| Total vSAN CPU Used (MHz)<br>• Summary \| Max vSAN CPU Ready<br>• Summary \| Worst VM Disk Latency |
| KPI | • KPI\|Max ESA Disk IOPS<br>• KPI\|Max ESA Disk Latency<br>• KPI\|Min Storage Pool Free Capacity<br>• KPI\|Sum Storage Pool Errors<br>• KPI\|Disk Groups Metrics<br>• KPI\|Sum Host VMKernel Packets Dropped<br>• KPI\|Count Disk Group Congestion Above 50<br>• KPI\|Max Disk Group Congestion<br>• KPI\|Sum Disk Group Errors<br>• KPI\|Min Disk Group Capacity Free<br>• KPI\|Min Disk Group Read Cache Hit Rate<br>• KPI\|Min Disk Group Write Buffer Free<br>• KPI\|Max Disk Group Read Cache/Write Buffer Latency<br>• KPI\|Max Capacity Disk Latency<br>• KPI \| Max Capacity Disk IOPS |
| IO Size | • vSAN \| Performance \| I/O Size (KB)<br>• vSAN \| Performance \| Read I/O Size (KB)<br>• vSAN \| Performance \| Write I/O Size (KB) |
| Resynchronization Status<br>( Metrics applicable for vSAN 6.7 and later) | • vSAN \| Resync \| Bytes left to resync (bytes)<br>• vSAN \| Resync \| Resyncing Objects |
| What If Analysis | This is an instanced metric<br>• vSAN \| What If \| Effective Free Space (GB) |
| Stretched Cluster | • vSAN\|Stretched Cluster\|Latency Between Sites\|Preferred and Secondary (ms)<br>• vSAN\|Stretched Cluster\|Latency Between Sites\|Preferred and Witness (ms)<br>• vSAN\|Stretched Cluster\|Latency Between Sites\|Secondary and Witness (ms) |
| File Share | • vSAN\|FileServices\|totalShareCount |
| File Service | • vSAN \| File Services \| File Shares Used Disk Space (GB)<br>• vSAN \| File Services \| Root FS Used Disk Space (GB)<br>• vSAN \| File Services \| File Shares Count |
| Slack Space | • vSAN\|Slack Space\|Internal Operations Capacity (GB)<br>• vSAN\|Slack Space\|Host Rebuild Capacity (GB)<br>• vSAN\|Slack Space\|Transient Capacity Used (GB) |

### Metrics for vSAN Enabled Host

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN enabled host.

Metrics for a vSAN enabled host include:

| Component | Metrics |
|---|---|
| Component Limit | • vSAN\|Component Limit\|Component Limit Used (%)<br>• vSAN\|Component Limit\|Total Component Limit<br>• vSAN\|Component Limit\|Used Component Limit |
| Disk Space | • vSAN\|Disk Space\|Disk Space Used (%)<br>• vSAN\|Disk Space\|Total Disk Space (GB)<br>• vSAN\|Disk Space\|Used Disk Space (GB) |
| Read Cache | • vSAN\|Read Cache\|Read Cache Reserved (%)<br>• vSAN\|Read Cache\|Reserved Read Cache Size (GB)<br>• vSAN\|Read Cache\|Total Read Cache Size (GB) |
| Performance Metrics | |
| • Network | • vSAN\|Performance\|Network\|Inbound Packets Loss Rate<br>• vSAN\|Performance\|Network\|Outbound Packets Loss Rate<br>• vSAN\|Performance\|Network\|<vnic>\|Inbound Packets Loss rate (%)<br>• vSAN\|Performance\|Network\|<vnic>\|Outbound Packets Loss Rate (%)<br>• vSAN\|Performance\|Network\|<vnic>\|Inbound Packets Per second<br>• vSAN\|Performance\|Network\|<vnic>\|Outbound Packets Per second<br>• vSAN\|Performance\|Network\|<vnic>\|Throughput Inbound (KBps)<br>• vSAN\|Performance\|Network\|<vnic>\|Throughput Outbound (KBps) |
| • CPU Utilization | • vSAN \| Performance \| CPU \| Ready (%)<br>• vSAN \| Performance \| CPU \| Usage (%)<br>• vSAN \| Performance \| CPU \| Used (MHz)<br>• vSAN \| Performance \| CPU \| Core Utilization (%) (For Hyper-Threading Technology) |
| • PCPU Utilization | • vSAN \| Performance \| PCPU \| Ready (%)<br>• vSAN \| Performance \| CPU \| PCPU Usage (%) |
| • Memory | • vSAN \| Performance \| Memory \| Usage (%)<br>• vSAN \| Performance \| Memory \| Used (GB) |

### Metrics for vSAN Datastore

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN datastore.

Datastore I/O metrics for vSAN datastore include:

- Datastore I/O|Reads Per Second (IOPS)
- Datastore I/O|Read Rate (KBps)
- Datastore I/O|Read Latency (ms)
- Datastore I/O|Writes Per Second (IOPS)
- Datastore I/O|Write Rate (KBps)

- Datastore I/O|Write Latency (ms)
- Datastore I/O|Outstanding I/O requests
- Datastore I/O|Congestion
- Capacity | Usable Capacity

### Metrics for vSAN Cache Disk

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN cache disk.

Metrics for vSAN cache disk include:

| Component | Metrics |
|---|---|
| Performance | • Performance\|Bus Resets<br>• Performance\|Commands Aborted Per Second<br>The following performance metrics are disabled by default:<br>• Performance\|Device Latency (ms)<br>• Performance\|Device Read Latency (ms)<br>• Performance\|Device Write Latency (ms)<br>• Performance\|Read Requests Per Second<br>• Performance\|Average Reads Per Second<br>• Performance\|Write Requests Per Second<br>• Performance\|Average Writes Per Second<br>• Performance\|Read Rate<br>• Performance\|Write Rate<br>• Performance\|Usage<br>• Performance\|HDD Errors |
| SCSI SMART Statistics | • SCSI SMART Statistics\|Health Status<br>• SCSI SMART Statistics\|Media Wearout Indicator<br>• SCSI SMART Statistics\|Write Error Count<br>• SCSI SMART Statistics\|Read Error Count<br>• SCSI SMART Statistics\|Power on Hours<br>• SCSI SMART Statistics\|Reallocated Sector Count<br>• SCSI SMART Statistics\|Raw Read Error Rate<br>• SCSI SMART Statistics\|Drive Temperature<br>• SCSI SMART Statistics\|Maximum Observed Drive Temperature<br>• SCSI SMART Statistics\|Drive Rated Max Temperature<br>• SCSI SMART Statistics\|Write Sectors TOT Count<br>• SCSI SMART Statistics\|Read Sectors TOT Count<br>• SCSI SMART Statistics\|Initial Bad Block Count<br>• SCSI SMART Statistics\|Worst Media Wearout Indicator<br>• SCSI SMART Statistics\|Worst Write Error Count<br>• SCSI SMART Statistics\|Worst Read Error Count<br>• SCSI SMART Statistics\|Worst Power-on Hours<br>• SCSI SMART Statistics\|Power Cycle Count<br>• SCSI SMART Statistics\|Worst Power Cycle Count |

*Table continued on next page*

*Continued from previous page*

| Component | Metrics |
|---|---|
| **NOTE** SMART data collection is disabled by default. To enable SMART data collection, ensure that the `Enable SMART data collection` instance identifier is set to true. For proper data collection, ensure that ESXi hosts in your vCenter Server inventory have CIM service enabled and CIM providers for each SMART metric installed. | • SCSI SMART Statistics\|Worst Reallocated Sector Count<br>• SCSI SMART Statistics\|Worst Raw Read Error Rate<br>• SCSI SMART Statistics\|Worst Driver Rated Max Temperature<br>• SCSI SMART Statistics\|Worst Write Sectors TOT Count<br>• SCSI SMART Statistics\|Worst Read Sectors TOT Count<br>• SCSI SMART Statistics\|Worst Initial Bad Block Count |
| Capacity | • vSAN\|Health\|Capacity\|Total Disk Capacity (GB)<br>• vSAN\|Health\|Capacity\|Used Disk Capacity (GB) |
| Congestion Health | • vSAN\|Health\|Congestion Health\|Congestion Value |
| Performance | • vSAN\|Performance\|Physical Layer Reads Per Second<br>• vSAN\|Performance\|Physical Layer Writes Per Second<br>• vSAN\|Performance\|Physical Layer Read Throughput (KBps)<br>• vSAN\|Performance\|Physical Layer Write Throughput (KBps)<br>• vSAN\|Performance\|Physical Layer Read Latency (ms)<br>• vSAN\|Performance\|Physical Layer Write Latency (ms)<br>• vSAN\|Performance\|Physical Layer Read Count<br>• vSAN\|Performance\|Physical Layer Write Count<br>• vSAN\|Performance\|Device Average Latency (ms)<br>• vSAN\|Performance\|Guest Average Latency (ms) |

### Metrics for vSAN Capacity Disk

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN capacity disk.

Metrics for vSAN capacity disk include:

| Component | Metrics |
|---|---|
| Performance | • Performance\|Bus Resets<br>• Performance\|Commands Aborted Per Second<br><br>The following performance metrics are disabled by default:<br>•<br>• Performance\|Device Latency (ms)<br>• Performance\|Device Read Latency (ms)<br>• Performance\|Device Write Latency (ms)<br>• Performance\|Read Requests Per Second<br>• Performance\|Average Reads Per Second<br>• Performance\|Write Requests Per Second<br>• Performance\|Average Writes Per Second<br>• Performance\|Read Rate<br>• Performance\|Write Rate<br>• Performance\|Usage<br>• Performance\|HDD Errors |
| SCSI SMART Statistics<br><br>**NOTE**<br>SMART data collection is disabled by default. To enable SMART data collection, ensure that the `Enable SMART data collection` instance identifier is set to true. For proper data collection, ensure that ESXi hosts in your vCenter Server inventory have CIM service enabled and CIM providers for each SMART metric installed. | • SCSI SMART Statistics\|Health Status<br>• SCSI SMART Statistics\|Media Wearout Indicator<br>• SCSI SMART Statistics\|Write Error Count<br>• SCSI SMART Statistics\|Read Error Count<br>• SCSI SMART Statistics\|Power on Hours<br>• SCSI SMART Statistics\|Reallocated Sector Count<br>• SCSI SMART Statistics\|Raw Read Error Rate<br>• SCSI SMART Statistics\|Drive Temperature<br>• SCSI SMART Statistics\|Maximum Observed Drive Temperature<br>• SCSI SMART Statistics\|Drive Rated Max Temperature<br>• SCSI SMART Statistics\|Write Sectors TOT Count<br>• SCSI SMART Statistics\|Read Sectors TOT Count<br>• SCSI SMART Statistics\|Initial Bad Block Count<br>• SCSI SMART Statistics\|Worst Media Wearout Indicator<br>• SCSI SMART Statistics\|Worst Write Error Count<br>• SCSI SMART Statistics\|Worst Read Error Count<br>• SCSI SMART Statistics\|Worst Power-on Hours<br>• SCSI SMART Statistics\|Power Cycle Count<br>• SCSI SMART Statistics\|Worst Power Cycle Count<br>• SCSI SMART Statistics\|Worst Reallocated Sector Count<br>• SCSI SMART Statistics\|Worst Raw Read Error Rate<br>• SCSI SMART Statistics\|Worst Driver Rated Max Temperature<br>• SCSI SMART Statistics\|Worst Write Sectors TOT Count<br>• SCSI SMART Statistics\|Worst Read Sectors TOT Count<br>• SCSI SMART Statistics\|Worst Initial Bad Block Count |
| Capacity | • vSAN\|Health\|Total Disk Capacity (GB)<br>• vSAN\|Health\|Used Disk Capacity (GB)<br>• vSAN\|FileServices\|FileSharesUsedDiskSpace<br>• vSAN\|FileServices\|RootFsUsedDiskSpace |

*Table continued on next page*

*Continued from previous page*

| Component | Metrics |
|---|---|
| Congestion Health | vSAN\|Health\|Congestion Value |
| Performance | • vSAN\|Performance\|Physical Layer Reads Per Second<br>• vSAN\|Performance\|Physical Layer Writes Per Second<br>• vSAN\|Performance\|Physical Layer Read Throughput (KBps)<br>• vSAN\|Performance\|Physical Layer Write Throughput (KBps)<br>• vSAN\|Performance\|Physical Layer Read Latency (ms)<br>• vSAN\|Performance\|Physical Layer Write Latency (ms)<br>• vSAN\|Performance\|Physical Layer Read Count<br>• vSAN\|Performance\|Physical Layer Write Count<br>• vSAN\|Performance\|Device Average Latency (ms)<br>• vSAN\|Performance\|Guest Average Latency (ms)<br>• vSAN\|Performance\|vSAN Layer Reads Per Second<br>• vSAN\|Performance\|vSAN Layer Writes Per Second<br>• vSAN\|Performance\|vSAN Layer Read Latency (ms)<br>• vSAN\|Performance\|vSAN Layer Write Latency (ms)<br>• vSAN\|Performance\|vSAN Layer Read Count<br>• vSAN\|Performance\|vSAN Layer Write Count<br>• vSAN \| Performance \| vSAN Layer Total IOPS |

Properties for vSAN capacity disk include:

- Name
- Size
- Vendor
- Type
- Queue Depth

### Metrics for vSAN Fault Domain Resource Kind

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN stretched cluster with fault domain.

Metrics for vSAN fault domain resource kind includes:

- CPU
  - Demand
    - Demand (MHz)
    - Demand without overhead (MHz)
    - Overhead (MHz)
    - Reserved Capacity (MHz)
    - Total Capacity (MHz)
    - VM CPU Usage (MHz)
    - Workload (%)
- Disk Space
  - Demand
    - Workload (%)
- Memory
  - Contention (KB)
  - Demand
    - Host Usage (KB)
    - Machine Demand (KB)
    - Reserved Capacity (KB)
    - Total Capacity (KB)
    - Utilization (KB)
    - Workload (%)
- vSAN
  - Disk Space
    - Total Disk Space (GB)
    - Used Disk Space (GB)

**Metrics for vSAN World**

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN world.

Metrics for vSAN world include:

- Summary|Total Number of VMs
- Summary|Total Number of Hosts
- Summary|Total IOPS
- Summary|Total Latency
- Summary|Total Number of Clusters
- Summary|Total Number of DiskGroups
- Summary|Total Number of Cache Disks
- Summary|Total Number of Capacity Disks
- Summary|Total Number of Datastores
- Summary|Total vSAN Disk Capacity (TB)
- Summary|Total vSAN Disk Capacity Used (TB)
- Summary|Remaining Capacity (TB)
- Summary|Remaining Capacity (%)
- Summary|Total Savings by Deduplication and Compression (GB)

### Metrics for vSAN File Server

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN File Server.

Metrics for vSAN File Server

| Component | Metrics |
|---|---|
| File Server | • vSAN \| Disk Space File Shares Used Disk Space (GB)<br>• vSAN \| Summary \| File Shares Count |

### Metrics for vSAN File Share

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN File Share.

Metrics for vSAN File Share

| Component | Metrics |
|---|---|
| Disk Space | • vSAN \| Disk Space \| Used Disk Space (GB) |
| Read Performance | • vSAN \| Performance \| Read Throughput Requested (MBps)<br>• vSAN \| Performance \| Read Throughput Transferred (MBps)<br>• vSAN \| Performance \| Read IOPS<br>• vSAN \| Performance \| Read Latency (ms) |
| Write Performance | • vSAN \| Performance \| Write Throughput Requested (MBps)<br>• vSAN \| Performance \| Write Throughput Transferred (MBps)<br>• vSAN \| Performance \| Write IOPS<br>• vSAN \| Performance \| Write Latency (ms) |

### Capacity Model for vSAN Objects

The capacity model introduced in VMware Cloud Foundation Operations 6.7 now extends the support for vSAN objects like, vSAN cluster, Fault domains, and Cache/Capacity disks. The Capacity tab provides Time Remaining data for the selected vSAN cluster, Fault domain, Cache/Capacity Disk objects. The information is presented in a graphical format.

#### Where You Find the Capacity Tab

In the menu, click **Environment**, then select a group, custom data center, application, or inventory object. The Object details page appears. Click the **Capacity** tab.

The VMware Aria OperationsVMware Cloud Foundation Operations defines the capacity model for the following vSAN resource containers:

- vSAN Cluster
  - Disk Space
- vSAN Fault Domain
  - CPU
  - Memory
  - Disk Space

- vSAN Cache/Capacity Disk
  - Disk Space

**Understanding the Capacity Tab**

For the selected vSAN resource, the capacity tab lists the capacity used and Time Remaining until the associated CPU, memory, and disk space resources, respectively, run out.

- If you select the vSAN cluster, the capacity tab lists the capacity used and time remaining until the associated disk space runs out.
- If you select the vSAN Fault Domain, the capacity tab lists the capacity used and time remaining until the associated CPU, memory, and disk space resources run out.
- If you select the vSAN Cache/Capacity Disk Space, the capacity tab lists capacity used and time remaining until the associated disk space runs out.

The available graph depicts - for your choice of CPU, memory, or disk space - the amount of resource used, plotted against time. A line on the graph shows 100 percent usable capacity and a trend line projects how swiftly resource use is approaching 100 percent. The time line shows when the selected resource is to reach capacity.

### Metrics for vSAN Storage Pool

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN storage pool.

| Component | Metrics |
|---|---|
| Disk I/O | • Disk I/O\|Physical Layer Read Latency (µs)<br>• Disk I/O\|Physical Layer Write Latency (µs)<br>• Disk I/O\|Physical Layer Read IOPS<br>• Disk I/O\|Physical Layer Write IOPS<br>• Disk I/O\|Physical Layer Read Throughput (KBps)<br>• Disk I/O\|Physical Layer Write Throughput (KBps)<br>• Disk I/O\|Total Bus Resets<br>• Disk I/O\|Total IOPS Aborted<br>• Disk I/O\|Number of Errors |
| Disk Space | • Disk Space\|Capacity (bytes)<br>• Disk Space\|Usage (%)<br>• Disk Space\|Used (bytes)<br>• Disk Space\|Workload(%) |

### Metrics for vSAN ESA Disk

The VMware Aria OperationsVMware Cloud Foundation Operations collects the metrics you use to monitor the performance of your vSAN ESA disk.

| Component | Metrics |
|---|---|
| vSAN Health\|Capacity | • vSAN\|Health\|Capacity\|Total Disk Capacity (GB)<br>• vSAN\|Health\|Capacity\|Used Disk Capacity (GB)<br>• vSAN\|Health\|Capacity\|Used Disk Capacity (%) |

*Table continued on next page*

*Continued from previous page*

| Component | Metrics |
|---|---|
| | • Disk I/O\|Physical Layer Write IOPS<br>• Disk I/O\|Physical Layer Read Throughput (KBps)<br>• Disk I/O\|Physical Layer Write Throughput (KBps)<br>• Disk I/O\|Total Bus Resets<br>• Disk I/O\|Total IOPS Aborted<br>• Disk I/O\|Number of Errors |
| vSAN Health\|Congestion | vSAN Health\|Congestion Health\|Congestion Value |
| vSAN Performance | • vSAN\|Performance\|Physical Layer Read IOPS<br>• vSAN\|Performance\|Physical Layer Write IOPS<br>• vSAN\|Performance\|Physical Layer Read Throughput (KBps)<br>• vSAN\|Performance\|Physical Layer Write Throughput (KBps)<br>• vSAN\|Performance\|Physical Layer Read Latency (µs)<br>• vSAN\|Performance\|Physical Layer Write Latency (µs)<br>• vSAN\|Performance\|Physical Layer Read Count<br>• vSAN\|Performance\|Physical Layer Write Count<br>• vSAN\|Performance\|Device Latency (µs)<br>• vSAN\|Performance\|Guest Latency (µs)<br>• vSAN\|Performance\|vSAN Layer Read IOPS<br>• vSAN\|Performance\|vSAN Layer Write IOPS<br>• vSAN\|Performance\|vSAN Layer Total IOPS<br>• vSAN\|Performance\|vSAN Layer Average Read Latency (µs)<br>• vSAN\|Performance\|vSAN Layer Average Write Latency (µs)<br>• vSAN\|Performance\|vSAN Layer Read Throughput (KBps)<br>• vSAN\|Performance\|vSAN Layer Write Throughput (KBps) |

### Google Cloud VMware Engine Metrics

VMware Aria Operations collects metrics for Google Cloud VMware Engine objects.

| Resource Kind | Metric Key | Metric Name | Metric | Description | Unit |
|---|---|---|---|---|---|
| GCVEAdapterInstance | configuration\|nodeCount | Configuration\|ESXi Hosts Count | Integer | The number of hosts in the project. | NA |
| | cost\|conversionFactor | Cost\|Conversion Factor | Double | Conversion factor for converting Google Cloud VMware Engine expenses from bill to currency configured on VMware Aria Operations. | NA |
| | cost\|total_aggregated_cost | Cost\|Total Aggregated Cost | Double | Total Aggregated Cost | CurrencyMonth |

*Table continued on next page*

*Continued from previous page*

| Resource Kind | Metric Key | Metric Name | Metric | Description | Unit |
|---|---|---|---|---|---|
| GCVEWorld | summary\| total_number_vcenters | Summary\|Total Number of vCenter Server(s) | Integer | Total number of vCenter Server(s). | NA |
| | cost\| total_aggregated_cost | Cost\|Total Aggregated Cost | Double | Total Aggregated Cost of all the GCVE Projects | CurrencyMonth |
| | cost\| total_reclaimable_cost | Cost\|Total Reclaimable Cost | Double | Cost associated with reclaimable resources within the GCVEWorld object. | CurrencyMonth |
| | cost\| capacity_remaining\|compute | Cost\|Capacity Remaining\| Compute Cost | Double | Cost of compute remaining. | CurrencyMonth |
| | cost\| capacity_remaining\|storage | Cost\|Capacity Remaining\| Storage Cost | Double | Cost of storage remaining. | CurrencyMonth |
| | cost\| capacity_used\| storage | Cost\|Capacity Used\|Storage Cost | Double | Cost of storage used. | CurrencyMonth |
| | cost\| capacity_used\| compute | Cost\|Capacity Used\|Storage Cost | Double | Cost of compute used. | CurrencyMonth |
| | cost\| potential_savings\| idle_vms | Cost\|Potential Savings\|Idle VMs Cost | Double | Cost Avoidance/ Savings possible in the event of reclamation of idle VMs. | CurrencyMonth |
| | cost\| potential_savings\| poweredOff_vms | Cost\|Potential Savings\|Powered Off VMs Cost | Double | Cost Avoidance/ Savings possible in the event of reclamation of powered off VMs. | CurrencyMonth |
| | cost\| potential_savings\| vm_snapshots | Cost\|Potential Savings\|VM Snapshots Cost | Double | Cost Avoidance/ Savings possible in the event of reclamation of VM snaphots. | CurrencyMonth |
| | cost\| potential_savings\| orphaned_disks | Cost\|Potential Savings\| Orphaned Disks Cost | Double | Cost Avoidance/ Savings possible in the event of reclamation of orphaned disks. | CurrencyMonth |
| | cost\| potential_savings\| oversized_vms | Cost\|Potential Savings\|Cost Optimization Opportunities | Double | Cost Avoidance/ Savings possible in the event oversized VMs are rightsized. | Percent |

*Table continued on next page*

*Continued from previous page*

| Resource Kind | Metric Key | Metric Name | Metric | Description | Unit |
|---|---|---|---|---|---|
| | cost\|potential_savings\|total_reclaimable_host_cost | Cost\|Potential Savings\|Reclaimable Host Cost | Double | Cost Avoidance/ Savings possible in the event of reclamation of potentially reclaimable hosts. | CurrencyMonth |
| | cost\|potential_savings\|total_cost_of_ownership | Cost\|Potential Savings\|Total Cost of Ownership | Double | Projected optimized cost possible in the event of recommended optimizations like reclamation and rightsizing. | CurrencyMonth |
| | cost\|potential_increase\|undersized_vms | Cost\|Potential Increase\|Undersized VMs Cost | Double | Potential cost increase from undersized VMs. | CurrencyMonth |
| | cost\|realized_savings\|total_realized_savings | Cost\|Realized Savings\|Total Realized Savings | Double | Realized savings or cost actually avoided through the action of reclamation or rightsizing that leads to tangible benefit. | Currency |
| | cost\|realized_savings\|realized_idle_savings | Cost\|Realized Savings\|Idle Savings | Double | Total cost avoided (or saved) following the reclamation of Idle VMs across all data centers. | Currency |
| | cost\|realized_savings\|realized_poweredOff_savings | Cost\|Realized Savings\|Powered Off Savings | Double | Total cost avoided (or saved) following the reclamation of powered off VMs across all data centers. | Currency |
| | cost\|realized_savings\|realized_oversized_savings | Cost\|Realized Savings\|Oversized Savings | Double | Total cost avoided (or saved) following the reclamation of oversized VMs across all data centers. | Currency |
| | cost\|realized_savings\|realized_reclaimableHost_savings | Cost\|Realized Savings\|Reclaimable Host Savings | Double | Total cost avoided (or saved) following the reclamation of | Currency |

*Table continued on next page*

*Continued from previous page*

| Resource Kind | Metric Key | Metric Name | Metric | Description | Unit |
|---|---|---|---|---|---|
| | | | | hosts across all data centers. | |
| | cost\| realized_savings\| realized_orphanedDiskSpace_savings | Cost\|Realized Savings\| Orphaned Disk Space Savings | Double | Total cost avoided (or saved) following the reclamation of orphaned disk space across all data centers. | Currency |
| | cost\| realized_savings\| realized_snapshotSpace_savings | Cost\|Realized Savings\|Snapshot Space Savings | Double | Total cost avoided (or saved) following the reclamation of snapshot space across all data centers. | Currency |
| PrivateCloud | summary\| resourceStatus | Summary\| Resource Status | String | Activity Status - Active / Deleted | NA |
| | summary\| createTime | Summary\|Create Time | String | Creation time of the Private Cloud. | NA |
| | summary\| updateTime | Summary\|Update Time | String | Last update time of the Private Cloud. | NA |
| | summary\| description | Summary\| Description | String | User-provided description for the Private Cloud. | NA |
| | configuration\| Name | Configuration\| Name | String | Name of the Private Cloud. | NA |
| | configuration\| networkName | Configuration\| Network Name | String | The relative resource name of the consumer VPC network. | NA |
| | configuration\| zone | ConfigurationZone | String | Zone in which the Private Cloud is deployed. | NA |
| | configuration\| region | Configuration\| Region | String | The name of region in which Private Cloud is deployed. | NA |
| | configuration\| instanceType | Configuration\| Instance Type | String | Type of host deployed in the Private Cloud. | NA |
| | configuration\| nodeCount | Configuration\| ESXi Hosts Count | Integer | The number of hosts in the Private Cloud. | NA |
| | configuration_max\| max_esxi_hosts_ | Configuration Maximums\| Maximum ESXi | Double | Represents the maximum number of ESXi Hosts that | NA |

*Table continued on next page*

*Continued from previous page*

| Resource Kind | Metric Key | Metric Name | Metric | Description | Unit |
|---|---|---|---|---|---|
| | per_private_cloud \|limit | Hosts per Private Cloud\|Hard Limit | | can be deployed on the given Private Cloud. | |
| | configuration_max\| max_esxi_hosts_per_private_cloud \|provisioned | Configuration Maximums\| Maximum ESXi Hosts per Private Cloud\|Provisioned | Double | Represents the number of provisioned ESXi hosts for the given Private Cloud. | NA |
| | configuration_max\| max_esxi_hosts_per_private_cloud \|limit_used | Configuration Maximums\| Maximum ESXi Hosts per Private Cloud\|Limit Used | Double | Represents the percentage of ESXi hosts limit used for the given Private Cloud. | Percent |
| | configuration_max\| max_clusters_per_private_cloud\| limit | Configuration Maximums\| Maximum Clusters per Private Cloud\| Hard Limit | Double | Represents the maximum number of clusters that can be deployed for the given Private Cloud. | NA |
| | configuration_max\| max_clusters_per_private_cloud\| provisioned | Configuration Maximums\| Maximum Clusters per Private Cloud\| Provisioned | Double | Represents the number of provisioned clusters for the given Private Cloud. | NA |
| | configuration_max\| max_clusters_per_private_cloud\| limit_used | Configuration Maximums\| Maximum Clusters per Private Cloud\| Limit Used | Double | Represents the percentage of cluster limit used for the given Private Cloud | Percent |
| | cost\| total_aggregated_cost | Cost\|Total Aggregated Cost | Double | Total Aggregated Cost | CurrencyMonth |
| GCVEBill | cost\| outstanding_expense | Cost\|Outstanding Expense | Double | Outstanding expense. | NA |
| | cost\| monthly_commit_expense | Cost\|Monthly Commit Expense | Double | Monthly expense of commit hosts. | NA |
| | cost\| monthly_on_demand_expense | Cost\|Monthly OnDemand Expense | Double | Monthly expense of on-demand hosts. | NA |
| | cost\| monthly_total_expense | Cost\|Monthly Total Expense | Double | Monthly Total Expense. | NA |

*Table continued on next page*

*Continued from previous page*

| Resource Kind | Metric Key | Metric Name | Metric | Description | Unit |
|---|---|---|---|---|---|
| | configuration\|OrgId | Configuration\|Organization ID | String | Represents the CSP Organization ID. | NA |
| | configuration\|Currency | Configuration\|Currency | String | Represents the billing currency. | NA |
| | configuration\|StatementStartDate | Configuration\|Statement Bill Start Date | String | Represents the billing start date. | NA |
| | configuration\|StatementEndDate | Configuration\|Statement Bill End Date | String | Represents the billing end date. | NA |
| | summary\|ytd_commit_expense | Summary\|YTD Commit Expense | Double | Year-to-date expense of commit hosts. | NA |
| | summary\|ytd_on_demand_expense | Summary\|YTD OnDemand Expense | Double | Year-to-date expense of on-demand hosts. | NA |
| | summary\|ytd_total_expense | Summary\|YTD Total Expense | Double | Year-to-date total expense. | NA |
| Component | cost\|total_component_expense | Cost\|Total Component Expense | Double | Month-to-date expense of the given component considering the statement start date as start of the month. | NA |
| | cost\|bill_lineitem_expense | Cost\|Component Expense | Double | Component expense. | NA |
| | configuration\|component_start_date | Configuration\|Component Start Date | String | Component start date. | NA |
| | configuration\|component_end_date | Configuration\|Component End Date | String | Component end date. | NA |
| | configuration\|component_usage_type | Configuration\|Component Usage Type | String | Component Usage Type - Commit/On-Demand. | NA |
| | configuration\|component_sku_description | Configuration\|Component SKU Description | String | Component SKU Description | NA |
| | configuration\|subscription_status | Configuration\|Subscription Status | String | Indicates if the subscription is still active or inactive. | NA |
| | summary\|number_of_units | Summary\|Number of Units Used | Integer | Number of hosts purchased with the price | NA |

*Table continued on next page*

*Continued from previous page*

| Resource Kind | Metric Key | Metric Name | Metric | Description | Unit |
|---|---|---|---|---|---|
| | | | | published on the component object. | |
| RegionInstance | summary\|Name | Summary\|Name | String | Name of the region associated with the Private Cloud. | NA |

### Metrics in VMware Cloud on AWS

The VMware Cloud on AWS collects metrics for objects.

**Table 517: VMware Cloud on AWS Metrics**

| Object Type | Metric Key | Metric Value | Description |
|---|---|---|---|
| Bill | Cost\| Monthly Commit Expense | Double | Represents the total amount spent on the Commit purchases for a month. |
| | Cost\|Monthly OnDemand Expense | Double | Represents the total amount spent on the OnDemand purchases for a month. |
| | Cost\| Monthly Total Expense | Double | Represents the total amount spent on the OnDemand and Commit purchases for a month. |
| | Cost\| Outstanding Expense | Double | Represents the daily Outstanding expenses. |
| Component | Cost\| Component Expense | Double | Represents the amount spent for the purchases of Commit or OnDemand components for a month. |
| Org Object | Configuration Maximum \| Number of hosts per Organization \| Soft Limit | Double | Represents the number of hosts per organization. |
| | Configuration Maximum \| Number of hosts per Organization \| Provisioned | Double | |
| | Configuration Maximum \| Number of hosts per Organization \| Soft Limit % Used | Double | |
| | Configuration Maximum \| Public IP Addresses (Elastic IPs) \| Soft Limit | Double | Represents the maximum number of IP addresses per organization. |
| | Configuration Maximum \| Public IP Addresses (Elastic IPs) \| Provisioned | Double | |
| | Configuration Maximum \| Public IP Addresses (Elastic IPs) \| Soft Limit % Used | Double | |
| | Configuration Maximum \| Number of SDDCs per Organization \| Soft Limit | Double | Represents the maximum number of SDDCs per organization. |

*Table continued on next page*

*Continued from previous page*

| Object Type | Metric Key | Metric Value | Description |
|---|---|---|---|
| | Configuration Maximum \| Number of SDDCs per Organization \| Provisioned Limit | Double | |
| | Configuration Maximum \| Number of SDDCs per Organization \| Soft Limit % Used | Double | |
| SDDC | VMC Configuration Maximums \| Linked VPC Count \| Limit | Double | Represents the maximum number of linked AWS VPCs per SDDC. |
| | VMC Configuration Maximums \| Linked VPC Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Linked VPC Count \| Limit % Used | Double | |
| | Configuration Maximum \| Max clusters \| Soft Limit | Double | Represents the maximum number of vSphere clusters per SDDC. |
| | Configuration Maximum \| Max clusters \| Hard Limit | Double | |
| | Configuration Maximum \| Max clusters \| Provisioned | Double | |
| | Configuration Maximum \| Max clusters \| Soft Limit % Used | Double | |
| | Configuration Maximum \| Max clusters \| Hard Limit % Used | Double | |
| | Configuration Maximum \| Maximum hosts per SDDC \| Limit | Double | Represents the maximum number of ESXi hosts per SDDC. |
| | Configuration Maximum \| Maximum hosts per SDDC \| Provisioned | Double | |
| | Configuration Maximum \| Maximum hosts per SDDC \| Limit % Used | Double | |
| | Configuration Maximum \| Maximum VMs per SDDC \| Limit | Double | Represents the maximum number of virtual machines per SDDC. |
| | Configuration Maximum \| Maximum VMs per SDDC \| Provisioned | Double | |
| | Configuration Maximum \| Maximum VMs per SDDC \| Limit % Used | Double | |
| | VMC Configuration Maximums \| MGW Gateway Firewall Rule Count \| Limit | Double | Represents the maximum number of Management Gateway Firewall rules. |
| | VMC Configuration Maximums \| MGW Gateway Firewall Rule Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| MGW Gateway Firewall Rule Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| CGW Gateway Firewall Rule Count \| Limit | Double | Represents the maximum number of Compute Gateway Firewall rules. |
| | VMC Configuration Maximums \| CGW Gateway Firewall Rule Count \| Provisioned | Double | |

*Table continued on next page*

*Continued from previous page*

| Object Type | Metric Key | Metric Value | Description |
|---|---|---|---|
| | VMC Configuration Maximums \| CGW Gateway Firewall Rule Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| Direct Connect private VIF Connection Count \| Limit | Double | Represents the maximum number of private virtual interfaces attached to one SDDC. |
| | VMC Configuration Maximums \| Direct Connect private VIF Connection Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Direct Connect private VIF Connection Count \| Limit % Used | Double | |
| Cluster Compute Resource | Configuration Maximum \| Min hosts per cluster for full SLA \| Status | Double | Represents the minimum number of ESXi per vSphere cluster that must be supported at full SLA. |
| | Configuration Maximum \| Minimum hosts per cluster for full SLA \| Limit Violated | Double | |
| | Configuration Maximum \| Min hosts per cluster for no SLA \| Limit | Double | Represents the minimum number of ESXi hosts per vSphere cluster with no SLA. |
| | Configuration Maximum \| Min hosts per cluster for no SLA \| Limit Violated | Double | |
| | Configuration Maximum \| Max hosts per cluster (including stretched clusters) \| Limit | Double | Represents the maximum number of ESXi hosts per vSphere cluster. This limit applies to both single-AZ clusters and stretched clusters. |
| | Configuration Maximum \| Max hosts per cluster (including stretched clusters) \| Provisioned | Double | |
| | Configuration Maximum \| Max hosts per cluster (including stretched clusters) \| Limit % Used | Double | |
| Resource Pool | CPU \| vCPUs Allocated to all Consumers | Double | Represents the number of vCPUs allocated to the vCenter and NSX management appliances in a regular-sized SDDC. |
| | Memory \| Memory Allocated to all Consumers | Double | Represents the RAM allocated to the vCenter and NSX management appliances in a large and regular sized SDDC. |
| Host System | Configuration Maximum \| VMs per host \| Limit | Double | Represents the maximum number of VMs per host. |
| | Summary \| Total Number of VMs | Double | |
| | VMC Configuration Maximum \| VMs per host \| Limit % Used | Double | |
| Logical Router | VMC Configuration Maximums \| IPSec VPN Tunnel Count \| Limit | Double | Represents the maximum number of IPsec VPN tunnels created per SDDC. |
| | VMC Configuration Maximums \| IPSec VPN Tunnel Count \| Provisioned | Double | |

*Table continued on next page*

*Continued from previous page*

| Object Type | Metric Key | Metric Value | Description |
|---|---|---|---|
| | VMC Configuration Maximums \| IPSec VPN Tunnel Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| L2VPN Client Count \| Limit | Double | Represents the maximum number of sites connecting to L2 VPN server per SDDC. |
| | VMC Configuration Maximums \| L2VPN Client Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| L2VPN Client Count \| Limit % Used | Double | |
| Logical Switch | VMC Configuration Maximums \| Logical Segment Count \| Limit | Double | Represents the maximum number of logical segments per SDDC. |
| | VMC Configuration Maximums \| Logical Segment Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Logical Segment Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| Logical Ports Count \| Limit | Double | Represents the maximum number of ports on a logical segment. |
| | VMC Configuration Maximums \| Logical Ports Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Logical Ports Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| Extended Network Count \| Limit | Double | Represents the maximum number of logical segments extended from on-premises. |
| | VMC Configuration Maximums \| Extended Network Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Extended Network Count \| Limit % Used | Double | |
| Router Service (NAT Rules) | VMC Configuration Maximums \| NAT Rule Count \| Limit | Double | Represents the maximum number of Compute Gateway NAT rules. |
| | VMC Configuration Maximums \| NAT Rule Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| NAT Rule Count \| Limit % Used | Double | |
| Group | VMC Configuration Maximums \| Distributed Firewall Grouping Object Count \| Limit | Double | Represents the maximum number of grouping objects (security groups). |
| | VMC Configuration Maximums \| Distributed Firewall Grouping Object Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Distributed Firewall Grouping Object Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| IP Address Count \| Limit | Double | Represents the maximum number of IP addresses that can be included in an IP set. |
| | VMC Configuration Maximums \| IP Address Count \| Provisioned | Double | |

*Table continued on next page*

*Continued from previous page*

| Object Type | Metric Key | Metric Value | Description |
|---|---|---|---|
| | VMC Configuration Maximums \| IP Address Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| Distributed Firewall Rule Count \| Limit | Double | Represents the maximum number of distributed firewall rules per grouping object (security group). |
| | VMC Configuration Maximums \| Distributed Firewall Rule Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Distributed Firewall Rule Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| vm Count \| Limit | Double | Represents the maximum number of VMs per grouping object (security group). |
| | VMC Configuration Maximums \| vm Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| vm Count \| Limit % Used | Double | |
| Firewall Sections | VMC Configuration Maximums \| Distributed Firewall Section Count \| Limit | Double | Represents the maximum number of distributed firewall sections. |
| | VMC Configuration Maximums \| Distributed Firewall Section Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Distributed Firewall Section Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| Distributed Firewall Rule Count \| Limit | Double | Represents the maximum number of distributed firewall rules across all sections groups such as, Emergency Rules, Infrastructure Rules, and so on. |
| | VMC Configuration Maximums \| Distributed Firewall Rule Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Distributed Firewall Rule Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| Distributed (Group_Name) Firewall Rule Count \| Limit | Double | Represents the maximum number of distributed firewall rules per section group. |
| | VMC Configuration Maximums \| Distributed (Group_Name) Firewall Rule Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Distributed (Group_Name) Firewall Rule Count \| Limit % Used | Double | |
| | VMC Configuration Maximums \| Distributed (Group_Name) Firewall Section Count \| Limit | Double | Represents the maximum number of distributed firewall sections per section group, such as, Emergency Rules, Infrastructure Rules, and so on. |

*Table continued on next page*

*Continued from previous page*

| Object Type | Metric Key | Metric Value | Description |
|---|---|---|---|
| | VMC Configuration Maximums \| Distributed (Group_Name) Firewall Section Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Distributed (Group_Name) Firewall Section Count \| Limit % Used | Double | |
| Virtual Machine | VMC Configuration Maximums \| Security Tag Count \| Limit | Double | Represents the maximum number of security tags per VM. |
| | VMC Configuration Maximums \| Security Tag Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| Security Tag Count \| Limit % Used | Double | |
| Management Cluster | VMC Configuration Maximums \| IPFIX Collector Count \| Limit | Double | Represents the maximum number of IPFIX Collectors configured. |
| | VMC Configuration Maximums \| IPFIX Collector Count \| Provisioned | Double | |
| | VMC Configuration Maximums \| IPFIX Collector Count \| Limit % Used | Double | |
| Datastore | Configuration Maximum \| Maximum datastore capacity that can be utilized \| Limit | Double | Represents the maximum datastore capacity that can be utilized. You can use up to 75% of available datastore capacity. Usage beyond this point creates a non-compliant environment as described in Service Level Agreement for VMware Cloud on AWS. |
| | Configuration Maximum \| Datastore capacity requiring remediation plan \| Limit | Double | Represents the datastore capacity that requires a remediation plan. You must prepare a remediation plan when capacity utilization nears 70%. You can either add hosts to augment datastore capacity or reduce storage utilization. |

**Table 518: VMware Cloud on AWS Metrics Properties**

| Object Type | Property Name | Property Value | Description |
|---|---|---|---|
| Bill | Configuration\| Currency | String | Represents the currency unit set in the VMware Cloud on AWS account by the customer. |
| | Configuration\| OrgId | String | Represents the organization ID for the associated bill. |
| | Configuration\| Statement Bill Start Date | String | Represents the start date of the statement bill. |
| | Configuration\| Statement Bill End Date | String | Represents the end date of the statement bill. |

*Table continued on next page*

*Continued from previous page*

| Object Type | Property Name | Property Value | Description |
|---|---|---|---|
| | Summary\| YTD Commit Expense | Double | Represents the total amount spent on the Commit purchases for the current calendar year until the last generated statement bill. |
| | Summary\| YTD OnDemand Expense | Double | Represents the total amount spent on the OnDemand purchases for the current calendar year until the last generated statement bill. |
| | Summary\|YTD Total Expense | Double | Represents the total amount spent on the Commit and OnDemand purchases for the current calendar year until the last generated statement bill. |
| Component | Configuration\| Component Start Date | String | Represents the billing start date of the component purchase. |
| | Configuration\| Component End Date | String | Represents the billing end date of the component purchase. |
| | Configuration\| Component SKU Description | String | Represents the SKU of the component. |
| | Configuration\| Component Service Type | String | Represents the component service type. |
| | Configuration\| Component Usage Type | String | Represents the component usage type. |
| | Configuration\| Subscription Status | boolean | Represents whether a Commit is still available for use. |
| | Summary\| Number of Units Used | Integer | Represents the total number of components. |
| Org | Configuration\| Id | String | Represents the organization ID. |
| | Configuration\| Name | String | Represents the organization name. |

### Metrics in NSX Adapter

The NSX adapter collects metrics for objects within its plug-in.

**Table 519: Metrics in the NSX On-Premise**

| Resource | Metrics | Metric Keys |
|---|---|---|
| Management Cluster | System Capacity<br>• Max Supported Count<br>• Max Threshold Percentage<br>• Min Threshold Percentage<br>• Usage Count<br>• Usage Count Percentage<br>• Severity | System Capacity Keys<br>• System Capacity\|<Object_Kind>\| MaxSupportedCount<br>• System Capacity\|<Object_Kind>\| MaxThresholdPercentage<br>• System Capacity\|<Object_Kind>\| MinThresholdPercentage<br>• System Capacity\|<Object_Kind>\|UsageCount<br>• System Capacity\|<Object_Kind>\| UsageCountPercentage |

*Table continued on next page*

*Continued from previous page*

| Resource | Metrics | Metric Keys |
|---|---|---|
| | | • System Capacity\|<Object_Kind>\|Severity |
| Transport Node | • CPU\|<br>– CPU Cores<br>– DPDK CPU Cores<br>– DPDK CPU Core Average Usage<br>– DPDK CPU Core Highest Usage<br>– Non-DPDK CPU Core Average Usage<br>– Non-DPDK CPU Core Highest Usage<br>• Memory<br>– Total<br>– Used<br>– Cache<br>– Total Swap<br>– Used Swap | • CPU Metric Keys<br>– Cpu\|Cores<br>– Cpu\|DPDKCores<br>– Cpu\|AvgDpdkCpuCoreUsage<br>– Cpu\|HighDpdkCpuCoreUsage<br>– Cpu\|AvgNonDpdkCpuCoreUsage<br>– Cpu\|HighNonDpdkCpuCoreUsage<br>• Memory metric keys<br>– Memory\|Total<br>– Memory\|Used<br>– Memory\|Cache<br>– Memory\|Total Swap<br>– Memory\|Used Swap |
| | File Systems\|<FileSystemMount>\|Used | FileSystems\|Used |
| | Statistics\|Interface\|<InterfaceID><br>• Received Data (bytes)<br>• Received Packets dropped<br>• Received Packets errors<br>• Received Framing errors<br>• Received Packets<br>• Transmitted Data (bytes)<br>• Transmitted Packets dropped<br>• Transmitted Packets errors<br>• Transmitted carrier losses detected<br>• Transmitted Packets<br>• Transmitted Collisions detected | Statistics Metric Keys<br>• stats\|Interface\|RxData<br>• stats\|Interface\|RxDropped<br>• stats\|Interface\|RxErrors<br>• stats\|Interface\|RxFrame<br>• stats\|Interface\|RxPackets<br>• stats\|Interface\|TxData<br>• stats\|Interface\|TxDropped<br>• stats\|Interface\|TxErrors<br>• stats\|Interface\|TxCarrier<br>• stats\|Interface\|TxPackets<br>• stats\|Interface\|TxColls |
| Load Balancer Service | • CPU Usage(%)<br>• Memory Usage(%)<br>• Active Transport Nodes<br>• Standby Transport Nodes<br>• Sessions:<br>– L4Average<br>– L4Current<br>– L4Maximum<br>– L4Total<br>– L7Average<br>– L7Current<br>– L7Maximum<br>– L7Total | • CPU Usage<br>• Memory Usage<br>• Active Transport Nodes<br>• Standby Transport Nodes<br>• Sessions\|L4Average<br>• Sessions\|L4Current<br>• Sessions\|L4Maximum<br>• Sessions\|L4Total<br>• Sessions\|L7Average<br>• Sessions\|L7Current<br>• Sessions\|L7Maximum<br>• Sessions\|L7Total |
| Load Balancer Virtual Server | • Statistics<br>– Bytes\|Inbound Bytes Total<br>– Bytes\|Average Inbound Bytes Per Second | • Statistics metric keys<br>– stats\|Bytes\|Inbound<br>– stats\|Bytes\|InboundRate |

*Table continued on next page*

*Continued from previous page*

| Resource | Metrics | Metric Keys |
|---|---|---|
| | – Bytes\|Outbound Bytes Total<br>– Bytes\|Average Outbound Bytes Per Second<br>– Http\|Http Request Rate<br>– Http\|Http Requests<br>– Packets\|Inbound Packets Total<br>– Packets\|Inbound Packets Rate<br>– Packets\|Outbound Packets Total<br>– Packets\|Outbound Packets Rate<br>– Packets\|Dropped<br>• Sessions<br>– Average Current Sessions Per Second<br>– Current Sessions<br>– Maximum Sessions<br>– Dropped Sessions<br>– Total Sessions | – stats\|Bytes\|Outbound<br>– stats\|Bytes\|OutboundRate<br>– stats\|Http\|RequestRate<br>– stats\|Http\|Requests<br>– stats\|Packets\|Inbound<br>– stats\|Packets\|InboundRate<br>– stats\|Packets\|Outbound<br>– stats\|Packets\|OutboundRate<br>– stats\|Packets\|Dropped<br>• Sessions metric keys<br>• – Sessions\|CurrentRate<br>– Sessions\|Current<br>– Sessions\|Maximum<br>– Sessions\|Dropped<br>– Sessions\|Total |
| Load Balancer Pool | • Statistics<br>– Bytes\|Inbound Bytes Total<br>– Bytes\|Average Inbound Bytes Per Second<br>– Bytes\|Outbound Bytes Total<br>– Bytes\|Average Outbound Bytes Per Second<br>– Http\|Http Request Rate<br>– Http\|Http Requests<br>– Packets\|Inbound Packets Total<br>– Packets\|Inbound Packets Rate<br>– Packets\|Outbound Packets Total<br>– Packets\|Outbound Packets Rate<br>– Packets\|Dropped<br>• Sessions<br>– Average Current Sessions Per Second<br>– Current Sessions<br>– Maximum Sessions<br>– Dropped Sessions<br>– Total Sessions | • Statistics metric keys<br>– stats\|Bytes\|Inbound<br>– stats\|Bytes\|InboundRate<br>– stats\|Bytes\|Outbound<br>– stats\|Bytes\|OutboundRate<br>– stats\|Http\|RequestRate<br>– stats\|Http\|Requests<br>– stats\|Packets\|Inbound<br>– stats\|Packets\|InboundRate<br>– stats\|Packets\|Outbound<br>– stats\|Packets\|OutboundRate<br>– stats\|Packets\|Dropped<br>• Sessions metric keys<br>– Sessions\|CurrentRate<br>– Sessions\|Current<br>– Sessions\|Maximum<br>– Sessions\|Dropped<br>– Sessions\|Total |
| Management Services | • Service Monitor Process ID<br>• Service Monitor Runtime state<br>• Service Process ID<br>• Service Runtime State | • ServiceMonitorProcessId<br>• ServiceMonitorRuntimeState<br>• ServiceProcessIds<br>• ServiceRuntimeState |
| Logical Router | Statistics<br>• Received Data (bytes) | Statistics metric keys<br>• stats\|RxData |

*Table continued on next page*

*Continued from previous page*

| Resource | Metrics | Metric Keys |
|---|---|---|
| | • Received Packets dropped<br>• Received Packets<br>• Transmitted Data (bytes)<br>• Transmitted Packets dropped<br>• Transmitted Packets | • stats\|RxDropped<br>• stats\|RxPackets<br>• stats\|TxData<br>• stats\|TxDropped<br>• stats\|TxPackets |
| | Configuration Maximums<br>• Router Port Count<br>• ARP Entries Count<br>• Tier 1 Router Count<br>• Route Map Count<br>• Route Maps\|<RouteMapName:RouteMapId>\|Rule Count<br>• Prefix List Count<br>• IP Prefix Lists\|<IPPrefixListName:IPPrefixListId>\|Prefix List Entries Count | Configuration Maximums metric keys<br>• configMax\|routerPortCount<br>• configMax\|routerArpEntryCount<br>**NOTE**<br>Metric applicable for T1 router.<br>• configMax\|tier1RouterCount<br>• configMax\|routeMapCount<br>• configMax\|RouteMaps\|routeMapRuleCount<br>**NOTE**<br>Metric applicable for T0 router.<br>• configMax\|prefixListCount<br>• configMax\|IPPrefixLists\|prefixListEntriesCount<br>**NOTE**<br>Metric applicable for T0 and T1 router. |
| Logical Switch | Statistics<br>• Inbound Bytes Total<br>• Inbound Bytes Dropped<br>• Inbound Bytes Throughput<br>• Outbound Bytes Total<br>• Outbound Bytes Dropped<br>• Outbound Bytes Throughput<br>• Inbound Packets Total<br>• Inbound Packets Dropped<br>• Inbound Packets Throughput<br>• Outbound Packets Total<br>• Outbound Packets Dropped<br>• Outbound Packets Throughput | Metric keys<br>• stats\|IngressBytes<br>• stats\|IngressBytesDropped<br>• stats\|IngressBytesThroughput<br>• stats\|IngressPackets<br>• stats\|IngressPacketsDropped<br>• stats\|IngressPacketsThroughput<br>• stats\|EgressBytes<br>• stats\|EgressBytesDropped<br>• stats\|EgressBytesThroughput<br>• stats\|EgressPackets<br>• stats\|EgressPacketsDropped<br>• stats\|EgressPacketsThroughput |
| Logical Switch Group | Configuration Maximums<br>• Logical Segment Count | Metric keys<br>• configMax\|LogicalSegmentCount |
| Management Appliances | Management Node Count | Management node count |
| Manager Node | • File Systems\|<FileSystemMount><br>– File System Id<br>– File System Type<br>– Total (KB)<br>– Used(KB)<br>– Used(%) | • File Systems Metric Keys<br>– FileSystems\|<FileSystemMount>\|FileSystemId<br>– FileSystems\|<FileSystemMount>\|Type<br>– FileSystems\|<FileSystemMount>\|Total<br>– FileSystems\|<FileSystemMount>\|Used |

*Table continued on next page*

*Continued from previous page*

| Resource | Metrics | Metric Keys |
|---|---|---|
| | | – FileSystems\|<FileSystemMount>\| usedPercentage |
| | Network Interfaces\|<InterfaceID>\|<br>• Received Data\|Bits per second<br>• Received Data\|Cumulative(bytes)<br>• Received Framing Errors\|Cumulative<br>• Received Framing Errors\|Per second<br>• Received Packets\|Cumulative<br>• Received Packets\|Per Second<br>• Received Packets Dropped\|Cumulative<br>• Received Packets Dropped\|Per second<br>• Received Packets Error\|Cumulative<br>• Received Packets Error\|Per second<br>• Transmitted Carrier losses detected\| Cumulative<br>• Transmitted Carrier losses detected\|Per second<br>• Transmitted Collisions detected\| Cumulative<br>• Transmitted Collisions detected\|Per second<br>• Transmitted Data\|Bits per second<br>• Transmitted Data\|Cumulative(bytes)<br>• Transmitted Packets\|Cumulative<br>• Transmitted Packets\|Per second<br>• Transmitted Packets Dropped\| Cumulative<br>• Transmitted Packets Dropped\|Per second<br>• Transmitted Packets errors\|Cumulative<br>• Transmitted Packets errors\|Per second | Network Interface metric keys<br>• Interfaces\|<InterfaceID>\|RxData\|BitsPerSecond<br>• Interfaces\|<InterfaceID>\|RxData\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxFrame\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxFrame\|PerSecond<br>• Interfaces\|<InterfaceID>\|RxPackets\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxPackets\|PerSecond<br>• Interfaces\|<InterfaceID>\|RxDropped\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxDropped\|PerSecond<br>• Interfaces\|<InterfaceID>\|RxErrors\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxErrors\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxCarrier\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxCarrier\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxColls\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxColls\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxData\|BitsPerSecond<br>• Interfaces\|<InterfaceID>\|TxData\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxPackets\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxPackets\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxDropped\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxDropped\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxErrors\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxErrors\|PerSecond |
| | CPU<br>• CPU Cores<br>• DPDK CPU Cores<br>• DPDK CPU Core Average Usage<br>• DPDK CPU Core Highest Usage<br>• Non-DPDK CPU Core Average Usage<br>• Non-DPDK CPU Core Highest Usage | CPU Metric Keys<br>• Cpu\|Cores<br>• Cpu\|DPDKCores<br>• Cpu\|AvgDpdkCpuCoreUsage<br>• Cpu\|HighDpdkCpuCoreUsage<br>• Cpu\|AvgNonDpdkCpuCoreUsage<br>• Cpu\|HighNonDpdkCpuCoreUsage |
| | Memory<br>• Total<br>• Used<br>• Cache<br>• Total Swap<br>• Used Swap | Memory metric keys<br>• Memory\|Total<br>• Memory\|Used<br>• Memory\|Cache<br>• Memory\|TotalSwap<br>• Memory\|UsedSwap |

*Table continued on next page*

*Continued from previous page*

| Resource | Metrics | Metric Keys |
|---|---|---|
| Controller Cluster | • Controller Node Count<br>• Cluster Status\|Controller Cluster Status<br>• Cluster Status\|Management cluster Status | Controller cluster metrics keys<br>• Cluster Status\|Controller Node Count<br>• ClusterStatus\|ControllerClusterStatus<br>• ClusterStatus\|ManagementClusterStatus<br>**NOTE**<br>These metrics are not collected for NSX version above 2.4 |
| Controller Node | • Connectivity Status\|Cluster Connectivity<br>• Connectivity Status\|Manager Connectivity<br>• File System ID<br>• File System Type<br>• Total(KB)<br>• Used(KB)<br>• Used(%)<br>• Network Interfaces\|<InterfaceID>\|<br>• Received Data\|Bits per second<br>• Received Data\|Cumulative(bytes)<br>• Received Framing Errors\|Cumulative<br>• Received Framing Errors\|Per second<br>• Received Packets\|Cumulative<br>• Received Packets\|Per Second<br>• Received Packets Dropped\|Cumulative<br>• Received Packets Dropped\|Per second<br>• Received Packets Error\|Cumulative<br>• Received Packets Error\|Per second<br>• Transmitted Carrier losses detected\| Cumulative<br>• Transmitted Carrier losses detected\|Per second<br>• Transmitted Collisions detected\| Cumulative<br>• Transmitted Collisions detected\|Per second<br>• Transmitted Data\|Bits per second<br>• Transmitted Data\|Cumulative(bytes)<br>• Transmitted Packets\|Cumulative<br>• Transmitted Packets\|Per second<br>• Transmitted Packets Dropped\| Cumulative<br>• Transmitted Packets Dropped\|Per second<br>• Transmitted Packets errors\|Cumulative<br>• Transmitted Packets errors\|Per second | **NOTE**<br>These metrics are not collected for NSX version above 2.4<br>• ConnectivityStatus\|ClusterConnectivity<br>• ConnectivityStatus\|ManagerConnectivity<br>• FileSystems\|<FileSystemMount>\|FileSystemId<br>• FileSystems\|<FileSystemMount>\|Type<br>• FileSystems\|<FileSystemMount>\|Total<br>• FileSystems\|<FileSystemMount>\|Used<br>• FileSystems\|<FileSystemMount>\|usedPercentage<br>• Interfaces\|<InterfaceID>\|RxData\|BitsPerSecond<br>• Interfaces\|<InterfaceID>\|RxData\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxFrame\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxFrame\|PerSecond<br>• Interfaces\|<InterfaceID>\|RxPackets\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxPackets\|PerSecond<br>• Interfaces\|<InterfaceID>\|RxDropped\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxDropped\|PerSecond<br>• Interfaces\|<InterfaceID>\|RxErrors\|Cumulative<br>• Interfaces\|<InterfaceID>\|RxErrors\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxCarrier\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxCarrier\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxColls\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxColls\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxData\|BitsPerSecond<br>• Interfaces\|<InterfaceID>\|TxData\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxPackets\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxPackets\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxDropped\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxDropped\|PerSecond<br>• Interfaces\|<InterfaceID>\|TxErrors\|Cumulative<br>• Interfaces\|<InterfaceID>\|TxErrors\|PerSecond |

*Table continued on next page*

*Continued from previous page*

| Resource | Metrics | Metric Keys |
|---|---|---|
| Router Service | RouterService | • BGP Neighbor:<BGPNeighborName>\|Service Router State\|Connection State<br>• BGP Neighbor:<BGPNeighborName>\|Advertised Routes\|Transport Nodes:<TransportNodeIP>\| Advertised Route Count<br>• BGP Neighbor:<BGPNeighborName>\|Advertised Routes\|Transport Nodes:<TransportNodeIP>Routes\|ASPath<br>• BGP Neighbor:<BGPNeighborName>\|Advertised Routes\|Transport Nodes:<TransportNodeIP>Routes\|Next Hop |

**Table 520: Metrics in the NSX on VMware Cloud on AWS**

| Resource | Metrics | Metric Keys |
|---|---|---|
| Logical Router | The following metrics are specify to Tier 0 Router.<br>Statistics \| Interface<br>• Received Data (Bytes)<br>• Received Packets<br>• Received Packets Dropped<br>• Transmitted Data<br>• Transmitted Received Data (Bytes)<br>• Transmitted Received Packets<br>• Transmitted Received Packets Dropped | Stats Metrics<br>Statistics \| Interface<br>• stats\|Interface\|RxData<br>• stats\|Interface\|RxPackets<br>• stats\|Interface\|RxDropped<br>• stats\|Interface\|TxData<br>• stats\|Interface\|TxPackets<br>• stats\|Interface\|TxDropped<br>**NOTE**<br>These metrics are only for Tier 0 Router. |
| Firewall Section Group | Configuration Maximums<br>• Distributed Firewall Section Count<br>• Distributed Firewall Rule Count<br>• MGW Gateway Firewall Rule Count<br>• CGW Gateway Firewall Rule Count<br>• Distributed Application Firewall Rule Count<br>• Distributed Application Firewall Section Count<br>• Distributed Environment Firewall Rule Count<br>• Distributed Environment Firewall Section Count | Configuration metric keys<br>• configMax\|MaxDistributedFirewallSections<br>• configMax\|MaxDistributedFirewallRules<br>• configMax\|MaxMGWGatewayFirewallRules<br>• configMax\|MaxCGWGatewayFirewallRules<br>• configMax\|MaxDistributedApplicationFirewallRules<br>• configMax\| MaxDistributedApplicationFirewallSections<br>• configMax\|MaxDistributedEnvironmentFirewallRules<br>• configMax\| MaxDistributedEnvironmentFirewallSections<br>• configMax\| MaxDistributedInfrastructureFirewallRules<br>• configMax\| MaxDistributedInfrastructureFirewallSections<br>• configMax\|MaxDistributedEmergencyFirewallRules |

*Table continued on next page*

*Continued from previous page*

| Resource | Metrics | Metric Keys |
|---|---|---|
| | • Distributed Infrastructure Firewall Rule Count<br>• Distributed Infrastructure Firewall Section Count<br>• Distributed Emergency Firewall Rule Count<br>• Distributed Emergency Firewall Section Count<br>• Distributed Ethernet Firewall Rule Count<br>• Distributed Ethernet Firewall Section Count<br>**NOTE**<br>These metrics are only for NSX on VMware Cloud on AWS. For NSX on-premise, the values for these metrics show zero. | • configMax\|MaxDistributedEmergencyFirewallSections<br>• configMax\|MaxDistributedEthernetFirewallRules<br>• configMax\|MaxDistributedEthernetFirewallSections<br>**NOTE**<br>These metrics are only for NSX on VMware Cloud on AWS. For NSX on-premise, the values for these metrics is shown as zero. |
| Logical Switch Group | Configuration Maximums<br>• Logical Segment Count<br>• Extended Network Count | Metric Keys<br>• configMax\|LogicalSegmentCount<br>• configMax\|ExtendedNetworkcount<br>**NOTE**<br>The metric (configMax\|ExtendedNetworkcount) is only for NSX on VMware Cloud on AWS. For NSX on-premise, its value is zero. |

**Sustainability Metrics**

Sustainability metrics are collected for virtual machine, host system, cluster compute resource, vSphere World, and Organization object types.

**Metric Names**

| Metric Name | Object Type | Description |
|---|---|---|
| Power\|Total Energy (Wh) | Virtual Machine | Total energy used.<br>Formula:<br>Total Energy (Wh) = Sum(Power\|Energy (Joule))/3600 |
| Power\|Total Energy (Wh) | Host System | Total energy used.<br>Formula: |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Object Type | Description |
|---|---|---|
| | | Total Energy (Wh) = Sum(Power\|Energy (Joule))/3600 |
| Sustainability\|CO2 Emission (kg) | Cluster Compute Resource | Carbon dioxide emissions. Calculated as power consumption* CO2 Emission rate<br>Formula:<br><br>CO2 Emission (kg) = Sum(Host System(Power\|Total Energy(Wh)))/1000 * 0.709 |
| Sustainability\|CO2 Emission before Virtualization (kg) | Cluster Compute Resource | Carbon dioxide emissions before virtualization, assuming that power consumption per physical server is 100W, reflecting a low end hardware specification.<br>Formula:<br><br>CO2 Emission before Virtualization (kg) = Summary\|Number of Running VMs * 0.1 * 0.709 |
| Sustainability\|CO2 Emission by Idle VMs (kg) | Cluster Compute Resource | Total carbon dioxide emission from all idle VMs. Calculated as CO2 emission rate * Power consumed by Idle VMs, where the rate is set at cluster property / 1000.<br>Formula:<br><br>CO2 Emission by Idle VMs (kg) = Sum(VM(Power\|Total Energy (Wh)), If VM(Summary\|Reclaimable Idle = 1) * 0.709<br><br>OR<br><br>CO2 Emission by Idle VMs (kg) = Power Wasted by Idle VMs (Wh) * 0.709 |
| Sustainability\|Electricity Cost Savings | Cluster Compute Resource | Estimated cost savings by virtualizing workloads. Calculated from the difference between power consumption before virtualization and after virtualization. The electricity cost is defined at the cluster custom property.<br>Formula:<br><br>Electricity Cost Savings = (Summary\|Number of Running VMs * 0.1 - Sum(Host System(Power\|Total Energy(Wh)))/1000) * 0.108 |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Object Type | Description |
|---|---|---|
| | | OR<br><br>Daily Electricity Cost Savings = (Summary\|Number of Running VMs * 0.1 - Power usage (kWh)) * 0.108 |
| Sustainability\|Power usage (kWh) | Cluster Compute Resource | Power usage calculated from all hosts in kWh.<br>Formula:<br><br>Power usage (kWh) = Sum(Host System(Power\|Total Energy(Wh))/ 1000 |
| Sustainability\|Power usage per GHz (Wh) | Cluster Compute Resource | Power usage efficiency. Calculated as power consumption over total GHz.<br>Formula:<br><br>Power usage per GHz (Wh) = Sum(Host System(Power\|Total Energy(Wh))/CPU\|Usage (MHz)/1000 |
| Sustainability\|Power Wasted by Idle VMs (Wh) | Cluster Compute Resource | Sum of electricity power used by all VMs classified as idle by the system.<br>Formula:<br><br>Power Wasted by Idle VMs (Wh) = Sum(VM(Power\|Total Energy (Wh)), If VM(Summary\|Reclaimable Idle = 1) |
| Sustainability\|Trees to Offset Idle VMs CO2 Emission | Cluster Compute Resource | Number of standard trees required to compensate CO2 emission of all Idle VMs. Based on 36.4 pound of carbon per tree.<br>Formula:<br><br>Trees to Offset Idle VMs CO2 Emission = Sum(VM(Power\|Total Energy (Wh)), If VM(Summary\| Reclaimable Idle = 1)/1000 * 0.709 / 16.511<br><br>OR<br><br>Trees to Offset Idle VMs CO2 Emission = Power Wasted by Idle VMs (Wh)/1000 * 0.709 / 16.511 |
| Sustainability\|CO2 Emission (kg) | vSphere World | Total carbon dioxide emissions. Calculated as sum of carbon emission from all clusters.<br>Formula:<br><br>CO2 Emission (kg) = Sum(Cluster(Sustainability\|CO2 Emission (kg)) |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Object Type | Description |
|---|---|---|
| Sustainability\|CO2 Emission before Virtualization (kg) | vSphere World | Total carbon dioxide emissions before virtualization. Calculated as the sum of CO2 emission from all clusters.<br>Formula:<br><br>CO2 Emission before Virtualization (kg) = Sum(Cluster(Sustainability\|CO2 Emission before Virtualization (kg))) |
| Sustainability\|CO2 Emission Avoided (t) | vSphere World | Carbon emissions avoided with virtualization. Calculated by difference in values of carbon emissions before and after virtualization and converting value in kg to Tonnes.<br>Formula:<br><br>CO2 Emission Avoided (t) = (CO2 Emission before Virtualization (kg) - CO2 Emission (kg)) / 1000 |
| Sustainability\|Electricity Cost Savings | vSphere World | Total estimated cost savings by virtualizing workloads at vSphere World. Calculated as the sum of Electricity cost savings.<br>Formula:<br><br>Electricity Cost Savings = Sum(Cluster(Daily Electricity Cost Savings)) |
| Sustainability\|Power Savings with Virtualization (%) | vSphere World | Percentage of power savings achieved by virtualization. Calculated by Formula = (Power usage before Virtualization - Power usage after Virtualization)/Power usage before Virtualization *100.<br>Formula:<br><br>Power Savings with Virtualization (%) =( Power usage Before Virtualization (kWh) - Power usage (kWh) ) / Power usage Before Virtualization (kWh) * 100 |
| Sustainability\|Power usage (kWh) | vSphere World | Power usage calculated from all hosts in kWh.<br>Formula:<br><br>Power usage (kWh) = Sum(Host System(Power\|Total Energy(Wh)))/ 1000 |
| Sustainability\|Power usage Before Virtualization (kWh) | vSphere World | Power usage assuming that each low range server consumes 0.1 kWh.<br>Formula: |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Object Type | Description |
|---|---|---|
| | | Power usage Before Virtualization (kWh) = Summary\|Number of Running VMs * 0.1 |
| Sustainability\|Power Wasted by Idle VMs (Wh) | vSphere World | Sum of electricity power used by all VMs classified as idle by the system.<br>Formula:<br><br>Power Wasted by Idle VMs (Wh) = Sum((VM(Power\|Total Energy (Wh)), If VM(Summary\|Reclaimable Idle = 1)) |
| Sustainability\|Adjusted Score\| Workload Efficiency (%) | Organization | Indicates efficiency based on resource wastage in the environment.<br>Formula:<br><br>Workload Efficiency (%) = 100 - (Wastage/Usage Ratio CPU + Wastage/Usage Ratio Memory + Wastage/Usage Ratio Disk) * 100 / 3 |
| Sustainability\|Adjusted Score\| Hardware Efficiency (%) | Organization | Indicates efficiency of hardware in the environment, based on the age of the hardware with the assumption that newer hardware is more energy efficient than older hardware.<br>Formula:<br><br>Hardware Efficiency (%) = ((10 - Average Age of Servers)/ 10 * 30%) + ((10 - Average Age of Storage)/ 10 * 30%) + ((10 - Average Age of Network)/ 10 * 20%) + ((10 - Average Age of Desktop)/ 10 * 20%)<br><br>The Average age of hardware variables in the formula is provided in the **Organization Details** page. |
| Sustainability\|Adjusted Score\| Resource Utilization (%) | Organization | Indicates utilization levels of the physical resources in the environment.<br>Formula:<br><br>Resource Utilization (%) = (Virtualized Server (%) * 44.45%) + (Virtualized Storage (%) * 33.33%) + (Virtualized Network (%) * 22.22%) |
| Sustainability\|Adjusted Score\| Virtualization (%) | Organization | Indicates the extent of virtualization in the environment.<br>Formula:<br><br>Virtualization (%) = (Server Virtualization (%) * 0.4) + (Storage Virtualization (%) * 0.3) + (Network |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Object Type | Description |
|---|---|---|
| | | Virtualization (%) * 0.2) + (Desktop Virtualization (%) * 0.1) |
| Sustainability\|Adjusted Score\|Power Source (%) | Organization | Indicates the sources of power being used in the environment. Usage of renewable energy sources leads to lesser carbon emission.<br>Formula:<br><br>Power Source Efficiency (%) = Sum(Adjusted Scores of each Power Source)<br><br>Adjusted Scores of Power Source = Power Source Green Factor * Power Source Share in Organization |
| Sustainability\|Hardware Utilization\| Server (%) | Organization | Indicates server utilization levels.<br>Formula:<br><br>Virtualized Server (%) = ((Total ESXi Hosts CPU Utilization (GHz) / Total ESXi Hosts CPU Capacity (GHz)) + (Total ESXi Hosts Memory Utilization (TB) / Total ESXi Hosts Memory Capacity (TB))) / 2 |
| Sustainability\|Hardware Utilization\| Storage (%) | Organization | Indicates storage utilization levels.<br>Formula:<br><br>Virtualized Storage (%) = (Total datastores utilization + Total RDM) / (Total datastores capacity + Total RDM + Physical Disk Overhead) |
| Sustainability\|Green Score (%) | Organization | A score that indicates how efficient or green your data center is.<br>Formula:<br><br>Green Score (%) = (Workload Efficiency (%) * 22.5%) + (Resource Utilization (%) * 12.5%) + (Virtualization Adoption (%) * 15%) + (Power Source Efficiency (%) * 37.5%) + (Hardware Power Efficiency (%) * 12.5%) |
| Sustainability\|Power Usage (kWh) | Organization | Sum of the power usage across all the clusters in the environment.<br>Formula:<br><br>Power Consumption = Sum ( (Cluster) Sustainability\|Power Usage (kWh) ) |
| Sustainability\|CO2 Emission (kg) | Organization | Total carbon emissions in the environment based on power consumed. |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Object Type | Description |
|---|---|---|
| | | Formula:<br><br>CO2 Emission = Power Consumption * CO2 Emission Ratio |
| Sustainability\|Workload Efficiency\|Idle VM Count | Organization | Number of Idle VMs, which are considered idle based on the Reclamation Idleness settings. |
| Sustainability\|Workload Efficiency\| Orphaned Disk Space (GB) | Organization | The aggregated size of all orphaned disks which are reclaimable based on the Orphaned Disk Reclamation settings. |
| Sustainability\|Workload Efficiency\| Oversized VM Count | Organization | Number of Oversized VMs, which are considered oversized by the capacity engine recommendation and are not in the Power Off state. |
| Sustainability\|Workload Efficiency\| Powered Off VM Count | Organization | The number of VMs which have been in the Powered Off state for the last 7 days. The number of days (7) can be configured from the Reclamation Settings. |
| Sustainability\|Workload Efficiency\| Snapshot Disk Space (GB) | Organization | The aggregated size of all snapshots which are reclaimable based on the Snapshot Reclamation settings. |

### Metrics for Synthetic Monitoring

These metrics display the average of the Synthetic Monitoring metrics at the Business Application and the Synthetic Monitoring endpoints.

### Synthetic Monitoring Endpoint Metrics

Synthetic Monitoring endpoint metrics are collected for Business Applications in which synthetic monitoring is enabled. These metrics are available for each endpoint.

**Synthetic Monitoring Endpoint Metrics**

| Metric Name | Metric Key |
|---|---|
| Content Transfer | Synthetic Monitoring:<api_id>\|Monitor:<geolocation>\| response_contentTransfer |
| DNS Lookup | Synthetic Monitoring:<api_id>\|Monitor:<geolocation>\| response_DNSLookup |
| Response Time (seconds) | Synthetic Monitoring:<api_id>\|Monitor:<geolocation>\| response_time |
| Server Processing (seconds) | Synthetic Monitoring:<api_id>\|Monitor:<geolocation>\| response_ServerProcessing |

*Table continued on next page*

*Continued from previous page*

| Metric Name | Metric Key |
|---|---|
| Status Code | Synthetic Monitoring:<api_id>\|Monitor:<geolocation>\|statusCode |
| TCP Connection (seconds) | Synthetic Monitoring:<api_id>\|Monitor:<geolocation>\|response_TCPConnection |
| TLS Handshake (seconds) | Synthetic Monitoring:<api_id>\|Monitor:<geolocation>\|response_TLSHandshake |

### Synthetic Monitoring Average Metrics

These metrics are an average of the Synthetic Monitoring metrics at the Business Application level.

**Synthetic Monitoring Average Metrics**

| Metric | Metric Key |
|---|---|
| Average Content Transfer (seconds) | Synthetic Monitoring\|average_content_transfer |
| Average DNS Lookup (seconds) | Synthetic Monitoring\|average_dns_lookup |
| Average Response Time (seconds) | Synthetic Monitoring\|average_response_time |
| Average Server Processing (seconds) | Synthetic Monitoring\|average_server_processing |
| Average TCP Connection | Synthetic Monitoring\|average_tcp_connection |
| Average TLS Handshake | Synthetic Monitoring\|average_tls_handshake |

### Metrics for Policies

Each policy is an object in VMware Aria OperationsVMware Cloud Foundation Operations. This allows you to track every change in the policy and create any dashboard, report, or alert on policy usage distribution.

VMware Aria OperationsVMware Cloud Foundation Operations collects metrics for the object type 'Policy'.

| Metric Key | Metric Name | Description |
|---|---|---|
| NumberOfEffectiveObjects | Number of Affected Objects | Represents the count of objects affected by the effective policy. |
| NumberOfVirtualMachines | Number of Affected vCenter Virtual Machines | Represents the number of vCenter Virtual Machines associated with a specific policy. |
| NumberOfCustomGroups | Number of Assigned Custom Groups | Represents the number of assigned custom groups for a specific policy. |
| NumberOfDefaultAssignedObjects | Number of Default Assigned Objects | Represents the number of objects that are assigned by default for a specific policy. This metric value is 0 for all policies except the default policy. |
| NumberOfDirectAssignedObjects | Number of Directly Assigned Objects | Represents the number of objects that are directly assigned for a specific policy. |

*Table continued on next page*

*Continued from previous page*

| Metric Key | Metric Name | Description |
|---|---|---|
| NumberOfInheritedObjects | Number of Objects Assigned By Scope | Represents the count of inherited objects for a specific policy. |

## Alert Definitions in VMware Aria OperationsVMware Cloud Foundation Operations

Alert definitions are a combination of symptoms and recommendations that identify problem areas in VMware Aria OperationsVMware Cloud Foundation Operations and generate alerts on which you act for those areas.

Alert definitions are provided for various objects in your environment. You can also create your own alert definitions.

Alert definitions are provided for various objects in your environment. You can also create your own alert definitions. See the *VMware Aria OperationsVMware Cloud Foundation Operations User Guide*.

### Cluster Compute Resource Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the Cluster Compute Resource objects in your environment.

**Health/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Symptom-based

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Fully-automated DRS-enabled cluster has CPU contention caused by less than half of the virtual machines. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster CPU contention at warning/ immediate/critical level<br>• > 0 descendant virtual machines have [ Virtual machine CPU demand at warning/ immediate/ critical level ]<br>• <= 50% of descendant virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level ]<br>• DRS Migration Threshold is not zero | 1. Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level.<br>2. Use the workload balance feature in VMware Cloud Foundation Operations to migrate one or more virtual machines to a different cluster.<br>3. Use vMotion to migrate some virtual machines to a different cluster if possible.<br>4. Add more hosts to the cluster to increase memory capacity.<br>5. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | within VMware Cloud Foundation Operations for right sizing of VMs. |
| Fully-automated DRS-enabled cluster has CPU contention caused by more than half of the virtual machines. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster CPU contention at warning/ immediate/critical level<br>• Cluster CPU demand at warning/ immediate/critical level<br>• > 50% of descendant virtual machines have [ Virtual machine CPU demand at warning/ immediate/critical level ]<br>• DRS Migration Threshold is not zero | 1. Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level.<br>2. Use the workload balance feature in VMware Cloud Foundation Operations to migrate one or more virtual machines to a different cluster.<br>3. Use vMotion to migrate some virtual machines to a different cluster if possible.<br>4. Add more hosts to the cluster to increase CPU capacity.<br>5. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for right sizing of VMs. |
| Fully-automated DRS-enabled cluster has CPU contention caused by overpopulation of virtual machines. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster CPU contention at warning/ immediate/critical level<br>• Cluster CPU workload at warning/ immediate/critical level<br>• = 0 descendant virtual machines have [ Virtual machine CPU demand at warning/ immediate/ critical level ]<br>• DRS Migration Threshold is not zero | 1. Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level.<br>2. Use the workload balance feature in VMware Cloud Foundation Operations to migrate one or more virtual machines to a different cluster.<br>3. Use vMotion to migrate some virtual machines to a different cluster if possible.<br>4. Add more hosts to the cluster to increase CPU capacity.<br>5. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for right sizing of VMs. |
| Fully-automated DRS-enabled cluster has high CPU workload. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster CPU workload above DT | 1. Check the applications running on the virtual machines in the cluster to determine whether high CPU workload is an expected behavior. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | • Cluster CPU workload at warning/ immediate/critical level | 2. Add more hosts to the cluster to increase CPU capacity.<br>3. Use vSphere vMotion to migrate some virtual machines to a different cluster if possible. |
| Fully-automated DRS-enabled cluster has memory contention caused by less than half of the virtual machines. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster memory contention at warning/immediate/critical level<br>• > 0 descendant virtual machines have [ Virtual machine memory workload at warning /immediate/ critical level ]<br>• <= 50% of descendant virtual machines have [Virtual machine memory workload at warning/ immediate/critical level ]<br>• DRS Migration Threshold is not zero | 1. Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level.<br>2. Use the workload balance feature in VMware Cloud Foundation Operations to migrate one or more virtual machines to a different cluster.<br>3. Use vMotion to migrate some virtual machines to a different cluster if possible.<br>4. To increase memory capacity add more hosts to the cluster.<br>5. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for right sizing of VMs. |
| Fully-automated DRS-enabled cluster has memory contention caused by more than half of the virtual machines. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster memory contention at warning/immediate/critical level<br>• Cluster memory workload at warning/immediate/critical level<br>• > 50% of descendant virtual machines have [ Virtual machine memory demand at warning/ immediate/critical level ]<br>• DRS Migration Threshold is not zero | 1. Check the migration threshold in the DRS settings for the cluster. Change it to a more aggressive level to enable DRS to balance the cluster workloads.<br>2. Use the workload balance feature in VMware Cloud Foundation Operations to migrate one or more virtual machines to a different cluster.<br>3. Use vMotion to migrate some virtual machines to a different cluster if possible.<br>4. Add more hosts to the cluster to increase memory capacity.<br>5. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for right sizing of VMs. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Fully-automated DRS-enabled cluster has memory contention caused by overpopulation of virtual machines. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster memory contention at warning/immediate/critical level<br>• Cluster memory workload at warning/immediate/critical level<br>• = 0 descendant virtual machines have [ Virtual machine memory demand at warning /immediate/ critical level ]<br>• DRS Migration Threshold is not zero | 1. Check the migration threshold in the DRS settings for the cluster. To enable DRS to balance the cluster workloads change it to a more aggressive level.<br>2. Use the workload balance feature in VMware Cloud Foundation Operations to migrate one or more virtual machines to a different cluster.<br>3. Use vMotion to migrate some virtual machines to a different cluster if possible.<br>4. Add more hosts to the cluster to increase memory capacity.<br>5. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for right sizing of VMs. |
| More than 5% of virtual machines in the cluster have memory contention due to memory compression, ballooning or swapping. | • Virtual machine memory limit is set AND<br>• > 5% of descendant virtual machines have [ virtual machine memory contention is at warning/ immediate/critical level] AND<br>• > 5% of descendant virtual machines have [ Virtual machine memory is compressed OR<br>• Virtual machine is using swap OR<br>• Virtual machine memory ballooning is at warning/immediate/critical level] | 1. Add more hosts to the cluster to increase memory capacity.<br>2. Use vMotion to migrate some virtual machines off the host or cluster. |
| Fully-automated DRS-enabled cluster has high memory workload and contention. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• Cluster memory contention above DT<br>• Cluster memory content is at warning/immediate/critical level<br>• Cluster memory workload at warning/immediate/critical level | 1. Check the applications running on the virtual machines in the cluster to determine whether high memory workload is an expected behavior.<br>2. Add more hosts to the cluster to increase memory capacity.<br>3. Use vSphere vMotion to migrate some virtual machines to a different cluster if possible. |
| vSphere High Availability (HA) failover resources are insufficient | vSphere High Availability (HA) failover resources are insufficient | To resolve this problem, use similar CPU and memory reservations for all virtual machines in the cluster. If this solution is not possible, consider using a different vSphere HA admission |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | control policy, such as reserving a percentage of cluster resource for failover. Alternatively, you can use advanced options to specify a cap for the slot size. For more information, see the vSphere Availability Guide. Hosts that have vSphere HA agent errors are not good candidates for providing failover capacity in the cluster and their resources are not considered for vSphere HA admission control purposes. If many hosts have a vSphere HA agent error, vCenter Server generates this event leading to the fault. To resolve vSphere HA agent errors, check the event logs for the hosts to determine the cause of the errors. After you resolve any configuration problems, reconfigure vSphere HA on the affected hosts or on the cluster. |
| vSphere HA master missing. | vCenter Server is unable to find a master vSphere HA agent (fault symptom) | |
| Proactive HA provider has reported health degradation on the underlying hosts. | Proactive HA provider reported host health degradation. | Contact your hardware vendor support. |

**Host System Alert Definitions**

The vCenter adapter provides alert definitions that generate alerts on the Host System objects in your environment.

**Health/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Standalone host has CPU contention caused by overpopulation of virtual machines.
Health

**Criticality**

Symptom-based

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Standalone host has CPU contention caused by less than half of the virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• Host CPU contention is at warning/ immediate/critical level | Use<br>1. Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | • > 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level]<br>• <= 50% of child virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] | when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Standalone host has CPU contention caused by more than half of the virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• Host CPU contention is at warning/ immediate/critical level<br>• Host CPU demand at warning/ immediate/critical level<br>• > 50% of child virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] | 1. Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Standalone host has CPU contention caused by overpopulation of virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• Host CPU contention is at warning/ immediate/critical level<br>• Host CPU demand at warning/ immediate/critical level<br>• = 0 child virtual machines have [Virtual machine CPU demand at warning/ immediate/critical level] | 1. Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Host in a cluster that does not have fully-automated DRS enabled has contention caused by less than half of the virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• [ DRS Enabled OR ! DRS fully automated ] | 1. Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | • Host CPU contention is at warning/ immediate/critical level<br>• > 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level]<br>• <= 50% of child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level] | resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Host in a cluster that does not have fully-automated DRS enabled has CPU contention caused by more than half of the virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• [ DRS Enabled OR ! DRS fully automated]<br>• Host CPU contention at warning/ immediate/critical level<br>• Host CPU demand at warning/ immediate/critical level<br>• > 50% of child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level] | 1. Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Host in a cluster that does not have fully-automated DRS enabled has CPU contention caused by overpopulation of virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• [ DRS Enabled OR ! DRS fully automated]<br>• Host CPU contention at warning/ immediate/critical level<br>• Host CPU demand at warning/ immediate/critical level<br>• = 0 child virtual machines have [Virtual machine CPU demand at warning /immediate/critical level | 1. Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Standalone host has memory contention caused by less than half of the virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• Host memory workload at warning/ immediate/critical level | 1. Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | • Host memory contention at warning/immediate/critical level<br>• > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/ critical level] | when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Upgrade the host to use a host that has larger memory capacity.<br>4. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Standalone host has memory contention caused by more than half of the virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• Host memory workload at warning/ immediate/critical level<br>• Host memory contention at warning/immediate/critical level<br>• > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/ critical level] | 1. Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Upgrade the host to use a host that has larger memory capacity.<br>4. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Standalone host has memory contention caused by overpopulation of virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• Host memory workload at warning/ immediate/critical level<br>• Host memory contention at warning/immediate/critical level<br>• = 0 child virtual machines have [Virtual machine memory workload at warning/ immediate/critical level] | 1. Add the host to a fully-automated-DRS cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Upgrade the host to use a host that has larger memory capacity.<br>4. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | Operations for recommended rightsizing of VMs. |
| Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by less than half of the virtual machines. | Symptoms include the following:<br>• [DRS Enabled OR ! DRS fully automated]<br>• Host memory contention at warning/immediate/critical level<br>• > 0 child virtual machines have [Virtual machine memory workload at warning/ immediate/critical level]<br>• <= 50% of child virtual machines have [Virtual machine memory workload at warning/ immediate/ critical level] | 1. Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by more than half of the virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• [DRS Enabled OR ! DRS fully automated]<br>• Host memory workload at warning/ immediate/critical level<br>• Host memory contention at warning/immediate/critical level<br>• > 50% of child virtual machines have [Virtual machine memory workload at warning /immediate/ critical level] | 1. Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Upgrade the host to use a host that has larger memory capacity.<br>4. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Host in a cluster that does not have fully-automated DRS enabled has memory contention caused by overpopulation of virtual machines. | Symptoms include the following:<br>• Host inside a cluster<br>• [DRS Enabled OR ! DRS fully automated]<br>• Host memory workload at warning/ immediate/critical level<br>• Host memory contention at warning/immediate/critical level<br>• = 0 child virtual machines have [Virtual machine memory workload at warning /immediate/critical level] | 1. Enable fully-automated DRS in the cluster to allow vSphere to move virtual machine as needed when resources are available on other hosts in the cluster.<br>2. Use vMotion to migrate some virtual machines with high CPU workload to other hosts that have available CPU capacity.<br>3. Upgrade the host to use a host that has larger memory capacity. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | 4. Right-size large virtual machines as it helps in reducing overall resource contention. Use the Reclaimable Capacity feature within VMware Cloud Foundation Operations for recommended rightsizing of VMs. |
| Host is experiencing high number of packets dropped. | Symptoms include the following:<br>• Host network received packets dropped<br>• Host network transmitted packets dropped | 1. Reduce the amount of network traffic being generated by virtual machines by moving some of them to a host with lower network traffic.<br>2. Verify the health of the physical network adapter, configuration, driver and firmware versions. |
| ESXi host has detected a link status 'flapping' on a physical NIC. | Physical NIC link state flapping (fault symptom). | ESXi disables the device to avoid the link flapping state. You might need to replace the physical NIC. The alert will be canceled when the NIC is repaired and functioning. If you replace the physical NIC, you might need to manually cancel the alert. |
| ESXi host has detected a link status down on a physical NIC. | Physical NIC link state down (fault symptom). | ESXi disables the device to avoid the link flapping state. You might need to replace the physical NIC. The alert will be canceled when the NIC is repaired and functioning. If you replace the physical NIC, you might need to manually cancel the alert. |
| Battery sensors are reporting problems. | Symptoms include the following:<br>• Battery sensor health is red OR<br>• Battery sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Baseboard Management Controller sensors are reporting problems. | Symptoms include the following:<br>• Baseboard Management Controller sensor health is red OR<br>• Baseboard Management Controller sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Fan sensors are reporting problems. | • Fan sensor health is red OR<br>• Fan sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Hardware sensors are reporting problems. | • Hardware sensor health is red OR<br>• Hardware sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Memory sensors are reporting problems. | • Memory sensor health is red OR<br>• Memory sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Path redundancy to storage device degraded | • A path to storage device went down<br>• Host has no redundancy to storage device | See KB topic, *Path redundancy to the storage device is degraded* (1009555) |
| Power sensors are reporting problems. | • Power sensor health is red OR<br>• Power sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Processor sensors are reporting problems. | • Processor sensor health is red<br>• Processor sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| IPMI System Event Log for the host is becoming full. | • SEL sensor health is red OR<br>• SEL sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Storage sensors are reporting problems. | • Storage sensor health is red OR<br>• Storage sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| System Board sensors are reporting problems. | • System board sensor health is red OR<br>• System board sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Temperature sensors are reporting problems. | • Temperature sensor health is red OR<br>• Temperature sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |
| Voltage sensors are reporting problems. | • Voltage sensor health is red OR<br>• Voltage sensor health is yellow | Change or replace the hardware if necessary. Contact the hardware vendor for assistance. After the problem is resolved, the alert will be canceled when the sensor that reported the problem indicates that the problem no longer exists. |

**Health/Critical**

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Critical

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Host has lost connection to vCenter. | Host disconnected from vCenter | Click "Open Host in vSphere Web Client" in the Actions menu at the top of Alert details page to connect to the vCenter managing this host and manually reconnect the host to vCenter Server. After the connection to the host is restored by vCenter Server, the alert will be canceled. |
| vSphere High Availability (HA) has detected a network-isolated host. | vSphere HA detected a network isolated host (fault symptom). | Resolve the networking problem that prevents the host from pinging its isolation addresses and communicating with other hosts. Make sure that the management networks that vSphere HA uses include redundancy. With redundancy, vSphere HA can communicate over more than one path, which reduces the chance of a host becoming isolated. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| vSphere High Availability (HA) has detected a possible host failure. | vSphere HA detected a host failure (fault symptom). | Find the computer that has the duplicate IP address and reconfigure it to have a different IP address. This fault is cleared and the alert canceled when the underlying problem is resolved, and the vSphere HA primary agent is able to connect to the HA agent on the host.<br>**NOTE**<br>You can use the Duplicate IP warning in the /`var/log/vmkernel` log file on an ESX host or the /var/log/messages log file on an ESXi host to identify the computer that has the duplicate IP address. |
| The host has lost connectivity to a dvPort. | Lost network connectivity to dvPorts (fault symptom). | Replace the physical adapter or reset the physical switch. The alert will be canceled when connectivity is restored to the dvPort. |
| The host has lost connectivity to the physical network. | Lost network connectivity (fault symptom). | To determine the actual failure or to eliminate possible problems, check the status of the vmnic in the vSphere Client or from the ESX service console:<br>• To check the status in the vSphere Client, select the ESX host, click **Configuration**, and then click **Networking**. The vmnics currently assigned to virtual switches appear in the diagrams. If a vmnic displays a red X, that link is currently down.<br>• From the service console, run the command:**`esxcfg-nics`**. The output that appears is similar to the following: Name PCI Driver Link Speed Duplex Description ---------- ---------- ---------- ---------- ---------- ---------- ---------- ---------- ---------- ---------- ---------- vmnic0 04:04.00 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet vmnic1 04:04.01 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet. The Link column shows the status of the link between the network adapter and the physical switch. The status can be either Up or Down. If some network adapters |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | are up and others are down, you might need to verify that the adapters are connected to the intended physical switch ports. To verify the connections, bring down each ESX host port on the physical switch, run **`esxcfg-nics -l"`**, and observe the affected vmnics.<br>Verify that the vmnic identified in the alert is still connected to the switch and configured properly:<br>• Make sure that the network cable is still connected to the switch and to the host.<br>• Make sure that the switch is connected to the system, is still functioning properly, and has not been inadvertently misconfigured. For more information, see the switch documentation.<br>• Check for activity between the physical switch and the vmnic. You can check activity by performing a network trace or observing activity LEDs.<br>• Check for network port settings on the physical switch.<br>To reconfigure the service console IP address if the affected vmnic is associated with a service console, see http://kb.vmware.com/kb/1000258 If the problem is caused by your hardware, contact your hardware vendor for replacement hardware. |
| The host lost connectivity to a Network File System (NFS) server. | Lost connection to NFS server (fault symptom). | 1. Verify the NFS server is running.<br>2. Check the network connection to make sure the ESX host can connect to the NFS server.<br>3. Determine whether the other hosts that use the same NFS mount are experiencing the same problem, and check the NFS server status and share points.<br>4. Make sure that you can reach the NFS server by logging into the service console and using **`vmkping`** to ping the NFS server: "vmkping <nfs server>". |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | 5. For advanced troubleshooting information, see http://kb.vmware.com/kb/1003967 |
| A fatal error occurred on a PCIe bus during system reboot. | A fatal PCIe error occurred. | Check and replace the PCIe device identified in the alert as the cause of the problem. Contact the vendor for assistance. |
| A fatal memory error was detected at system boot time. | A fatal memory error occurred. | Replace the faulty memory or contact the vendor. |

**Health/Immediate**

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Immediate

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| The host has lost redundant connectivity to a dvPort. | Lost network redundancy to DVPorts (fault symptom). | Replace the physical adapter or reset the physical switch. The alert will be canceled when connectivity is restored to the DVPort. |
| The host has lost redundant uplinks to the network. | Lost network redundancy (fault symptom). | To determine the actual failure or to eliminate possible problems, first connect to ESX through SSH or the console:<br>1. Identify the available uplinks by running **`esxcfg-nics -l`**.<br>2. Remove the reported vmnic from the port groups by running **`esxcfg-vswitch -U <affected vmnic#>`**; affected vSwitch.<br>3. Link available uplinks to the affected port groups by running **`esxcfg-vswitch -L <available vmnic#>`**; affected vSwitch.<br>Next, check the status of the vmnic in vSphere Client or the ESX service console:<br>1. In vSphere Client, select the ESX host, click **Configuration**, and then click **Networking**. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | | The vmnics currently assigned to virtual switches appear in the diagrams. If a vmnic displays a red X, that link is currently unavailable.<br>2. From the service console, run **`esxcfg-nics -l`**. The output that appears is similar to the following example: Name PCI Driver Link Speed Duplex Description.<br>---------- vmnic0 04:04.00 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet vmnic1 04:04.01 tg3 Up 1000Mbps Full Broadcom BCM5780 Gigabit Ethernet. The Link column shows the status of the link between the network adapter and the physical switch. The status can be either Up or Down. If some network adapters are up and others are down, you might need to verify that the adapters are connected to the intended physical switch ports. To verify the connections, shut down each ESX host port on the physical switch, run the "esxcfg-nics -l" command, and observe the affected vmnics. Verify that the vmnic identified in the alert is still connected to the switch and configured properly:<br>1. Make sure that the network cable is still connected to the switch and to the host.<br>2. Make sure that the switch is connected to the system, is still functioning properly, and was not inadvertently misconfigured. (See the switch documentation.)<br>3. Perform a network trace or observe activity LEDs to check for activity between the physical switch and the vmnic.<br>4. Check for network port settings on the physical switch.<br>If the problem is caused by hardware, contact your hardware vendor for a hardware replacement. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| A PCIe error occurred during system boot, but the error is recoverable. | A recoverable PCIe error occurred. | The PCIe error is recoverable, but the system behavior is dependent on how the error is handled by the OEM vendor's firmware. Contact the vendor for assistance. |
| A recoverable memory error has occurred on the host. | A recoverable memory error occurred. | Since recoverable memory errors are vendor-specific, contact the vendor for assistance. |

**Risk/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Risk

**Criticality**

Symptom-based

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| ESXi Host is violating vSphere 5.5 Hardening Guide. | • Active directory authentication disabled OR<br>• Non-compliant NTP service startup policy OR<br>• SSH service is running OR<br>• NTP service stopped OR<br>• Non-compliant timeout value for automatically disabling local and remote shell access OR<br>• vSphere Authentication Proxy not used for password protection when adding ESXi hosts to active directory OR<br>• Persistent logging disabled OR<br>• Bidirectional CHAP for iSCSI traffic disabled OR<br>• Non-compliant firewall setting to restrict access to NTP client OR<br>• NTP server for time synchronization not configured OR<br>• Non-compliant ESXi Shell service startup policy OR<br>• Non-compliant firewall setting to restrict access to SNMP server OR<br>• ESXi Shell service is running OR | Fix the vSphere 5.5 Hardening Guide Rules Violations according to the recommendations in the vSphere5 Hardening Guide |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | • Non-compliant DCUI service startup policy OR<br>• Dvfilter bind IP address configured OR<br>• Non-compliant SSH service startup policy OR<br>• DCUI service is running OR<br>• Non-compliant idle time before an interactive shell is automatically logged out OR<br>• Non-compliant DCUI access user list OR<br>• Remote syslog is not enabled | |

**VMware Aria Automation Alert Definitions**

Alert definitions are combinations of symptoms and recommendations that identify problem areas in your environment and generate alerts on which you can act.

Symptom and alert definitions are defined for VMware Aria Automation objects. The alerts are population-based alerts based on the risk or health of a certain percentage of child objects. There are no alerts generated for network profiles.

The health and risk thresholds are as follows:

**Health**

- When 25%-50% of the child objects have health issues, the parent object will trigger an alert with a Warning health level.
- When 50%-75% of the child objects have health issues, the parent object will trigger an alert with an Immediate health level.
- When 75%-100% of the child objects have health issues, the parent object will trigger an alert with a Critical health level.

**Risk**

- When 25%-50% of the child objects have risk issues, the parent object triggers an alert with a Warning risk level.
- When 50%-75% of the child objects have risk issues, the parent object triggers an alert with an Immediate risk level.
- When 75%-100% of the child objects have risk issues, the parent object triggers an alert with a Critical risk level.

**Cloud Zone Limits**

- When the Project approaches 70% of Memory allocation limit in one of the cloud zones, VMware Aria Automation triggers an alert.
- When the Project approaches 70% of Storage allocation limit in one of the cloud zones, VMware Aria Automation triggers an alert.
- When the Project approaches 70% of vCPU allocation limits in one of the cloud zones, VMware Aria Automation triggers an alert.

**Cloud Zone**

- Cloud zone has 60 days remaining until capacity runs out.

- Cloud Zone has less than 30 percent of capacity remaining.
- Cloud Zone has more than 20 percent of reclaimable capacity.

**Project**

- Project has more than 20 percent of reclaimable capacity.
- Project is approaching 70% of allocation limits.

### vSAN Alert Definitions

VMware Aria OperationsVMware Cloud Foundation Operations generates an alert if a problem occurs with the components in the storage area network that the vSAN adapter is monitoring.

#### Alerts for the vSAN Cluster Object

Alerts on the vSAN Cluster object have health, risk, and efficiency impact.

**Table 521: vSAN Cluster Object Health Alert Definitions**

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| Basic (unicast) connectivity check (normal ping) has failed on vSAN host. | Storage | Configuration | Triggered when basic (unicast) connectivity check (normal ping) has failed on the vSAN host due to network misconfiguration. |
| Check the free space on physical disks in the vSAN cluster. | Storage | Availability | Triggered when a check of free space on physical disks in the vSAN cluster results in an error or warning. |
| CLOMD process on the host has issues and impacting the functionality of vSAN cluster. | Storage | Availability | Triggered when CLOMD process on the host has issues and impacting the functionality of vSAN cluster. |
| Disk load variance between some vSAN disks exceeded the threshold value. | Storage | Performance | Triggered when disk load variance between some vSAN disks exceeded the threshold value.<br><br>vSAN cannot perform the load balance properly. |
| Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster. | Storage | Configuration | Host ESXi version and the vSAN disk format version is incompatible with the other hosts and disks in a vSAN cluster. |
| Host has invalid unicast agent and impacting the health of vSAN Stretched Cluster. | Storage | Configuration | Triggered when the host has invalid unicast agent and impacting the health of vSAN Stretched Cluster.<br><br>An invalid unicast agent on the host can cause a communication malfunction with the witness host. |
| Host in a vSAN cluster does not have a VMkernel NIC configured for vSAN traffic. | Network | Configuration | Triggered when the host in a vSAN cluster does not have a VMkernel NIC configured for vSAN traffic.<br><br>**NOTE** |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| | | | Even if an ESXi host is part of the vSAN cluster, but is not contributing storage, it must still have a VMkernel NIC configured for vSAN traffic. |
| Host in a vSAN cluster has connectivity issues and vCenter Server does not know its state. | Network | Configuration | Triggered when the host in a vSAN cluster has connectivity issues and vCenter Server does not know its state. |
| Host in a vSAN cluster has IP multicast connectivity issue. | Network | Configuration | Triggered when the host in a vSAN cluster has IP multicast connectivity issue. It means that multicast is most likely the root cause of a vSAN network partition. |
| Host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host. | Storage | Configuration | Triggered when the host is either running an outdated version of the vSAN Health Service VIB or It is not installed on the host. |
| Network latency check of vSAN hosts failed. It requires < 1 ms RTT. | Network | Configuration | Triggered if network latency check of vSAN hosts is greater than or equal to 1 ms RTT. |
| One or more hosts in the vSAN cluster have misconfigured multicast addresses. | Network | Configuration | Triggered when one or more hosts in the vSAN cluster have misconfigured multicast addresses. |
| One or more physical disks on vSAN host is experiencing software state health issues. | Storage | Availability | Triggered when one or more physical disks on vSAN host is experiencing software state health issues. |
| One or more vSAN enabled hosts are not in the same IP subnet. | Network | Configuration | Triggered when one or more vSAN enabled hosts are not in the same IP subnet. |
| Overall health of the physical disks in a vSAN Cluster is impacted. | Storage | Availability | Triggered when overall health of the physical disks in a vSAN Cluster is impacted. See the health status of each physical disk individually on all the hosts. |
| Overall health of VMs residing on vSAN datastore is reporting issues. | Storage | Availability | Triggered when overall health of the VMs on a vSAN datastore is impacted. |
| Overall health of vSAN objects is reporting issues. | Storage | Availability | Triggered when overall health of vSAN objects is reporting issues. |
| Ping test with large packet size between all VMKernel adapters with vMotion traffic enabled has issues. | Network | Configuration | Triggered when ping test with large packet size between all VMKernel adapter with vMotion traffic enabled is impacted. |
| Ping test with small packet size between all VMkernel adapters with vMotion traffic enabled has issues. | Network | Configuration | Triggered when ping test with small packet size between all VMKernel adapter with vMotion traffic enabled is impacted. |
| Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster. | Storage | Performance | Site latency between two fault domains and the witness host has exceeded the recommended threshold values in a vSAN Stretched cluster. |
| Statistics collection of vSAN performance service is not working correctly. | Storage | Availability | Triggered when statistics collection of vSAN performance service is not working correctly. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| | | | This means that statistics collection or writing statistics data to storage have failed for three consecutive intervals. |
| MTU check (ping with large packet size) has failed on vSAN host. | Storage | Configuration | Triggered when MTU check (ping with large packet size) has failed on vSAN environment due to some MTU misconfiguration in the vSAN network. |
| The preferred fault domain is not set for the witness host in a vSAN Stretched cluster. | Storage | Configuration | Triggered when the preferred fault domain is not set for the witness host in a vSAN Stretched cluster and affecting the operations of vSAN Stretched cluster. |
| Unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster. | Storage | Configuration | Triggered when unicast agent is not configured on the host and affecting operations of vSAN Stretched cluster. |
| vCenter Server has lost connection to a host that is part of a vSAN cluster. | Storage | Availability | Triggered when the host that is part of a vSAN cluster is in disconnected state or not responding and vCenter Server does not know its state. |
| vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster. | Storage | Configuration | Triggered when vSAN Cluster contains host whose ESXi version does not support vSAN Stretched Cluster. |
| vSAN cluster has issues in electing stats master of vSAN Performance service. This affects the functionality of vSAN Performance service. | Storage | Configuration | Triggered when vSAN cluster has issues in electing stats controller of vSAN Performance service. |
| vSAN cluster has multiple network partitions. | Network | Configuration | Triggered when vSAN cluster has multiple network partitions due to a network issue. |
| vSAN Cluster has multiple Stats DB objects which are creating conflicts and affecting vSAN Performance Service. | Storage | Configuration | Triggered when vSAN cluster has issues in electing stats controller of vSAN Performance service.<br><br>This affects the functionality of vSAN Performance service. |
| vSAN disk group has incorrect deduplication and compression configuration. | Storage | Configuration | Triggered when vSAN disk group has incorrect deduplication and compression configuration. |
| vSAN has encountered an issue while reading the metadata of a physical disk. | Storage | Availability | Triggered when vSAN has encountered an issue while reading the metadata of a physical disk and cannot use this disk. |
| vSAN health service is not installed on the host. | Storage | Configuration | Triggered when vSAN health service is not installed on the host. |
| vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster. | Storage | Configuration | Triggered when vSAN host and its disks have inconsistent deduplication and compression configuration with the cluster. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN is unable to retrieve the physical disk information from host. | Storage | Availability | Triggered when vSAN is unable to retrieve the physical disk information from host. vSAN Health Service may not be working properly on this host. |
| vSAN Performance Service is not enabled. | Storage | Configuration | Triggered when vSAN Performance Service is not enabled. |
| vSAN Performance Service is unable to communicate and retrieve statistics from host. | Storage | Configuration | Triggered when vSAN Performance Service is unable to communicate and retrieve statistics from host. |
| vSAN Performance Service network diagnostic mode is enabled for more than 24 hours. | Storage | Configuration | Triggered when the network diagnostic mode in vSAN Performance Service is enabled for more than 24 hours. |
| vSAN Stretched cluster contains a witness host without a valid disk group. | Storage | Configuration | Triggered when vSAN Stretched cluster contains a witness host without a valid disk group.<br><br>If the witness host does not have any disk claimed by vSAN then its fault domain is not available. |
| vSAN Stretched cluster does not contain a valid witness host. | Storage | Configuration | Triggered when vSAN Stretched cluster does not contain a valid witness host.<br><br>This affects the operations of vSAN Stretched cluster. |
| vSAN Stretched cluster does not contain two valid fault domains. | Storage | Configuration | Triggered when vSAN Stretched cluster does not contain two valid fault domains. |
| vSAN Stretched cluster has inconsistent configuration for Unicast agent. | Storage | Configuration | Triggered when vSAN Stretched cluster contains multiple unicast agents.<br><br>This means multiple unicast agents were set on non-witness hosts. |
| vSAN witness host has an invalid preferred fault domain. | Storage | Configuration | Triggered when vSAN witness host has an invalid preferred fault domain. |
| Witness host is a part of vSAN Stretched cluster. | Storage | Configuration | Triggered when witness host is a part of the vCenter cluster, which forms vSAN Stretched cluster. |
| Witness host resides in one of the data fault domains. | Storage | Configuration | Triggered when witness host resides in one of the data fault domains.<br><br>This affects the operations of vSAN Stretched cluster. |
| Witness appliance upgrade to vSphere 7.0 or higher with caution. | Storage | Configuration | Triggered when you want to upgrade the witness applicance to vSphere 7.0 or higher. |
| vSAN Support Insight is not enabled for the environment. | Storage | Configuration | Triggered when vSAN Support Insight is not enabled for the environment. |
| LSI 3108 controller’s advanced configuration values is different from recommended values. | Storage | Configuration | Triggered when the LSI-3108 based controller configuration values differs from vSAN configuration recommended values. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN Cluster flash read cache reservation is approaching capacity. | Application | Performance | Triggered when the flash read cache reservation in a vSAN cluster is less than 20%.<br><br>Cleared by adding more flash storage to the read-cache. |
| Some vSAN hosts are not compliant with the hyperconverged cluster configuration. | Storage | Configuration | Triggered when one of the host in vSAN cluster is not compliant with the hyperconverged cluster configuration. |
| Some vSAN hosts are not compliant for VMware vSphere Distributed Switch configuration. | Storage | Configuration | Triggered when one of the host in vSAN cluster is not compliant with the VMware vSphere Distributed Switch configuration. |
| Dual encryption is applied on virtual machines of a vSAN cluster. | Storage | Availability | Triggered when dual encryption is applied on a virtual machines of a vSAN cluster. |

**Table 522: vSAN Cluster Object Risk Alert Definitions**

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| After one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects | Storage | Capacity | Triggered when after one additional host failure, vSAN Cluster will not have enough resources to rebuild all objects. |
| Capacity disk used for vSAN is smaller than 255 GB (default max component size). | Storage | Performance | Triggered when a capacity disk used for vSAN is smaller than 255 GB (default max component size), so virtual machines that run on the vSAN Datastore might experience disk space issues. |
| Capacity disk used for vSAN is smaller than 255 GB (default max component size). | Storage | Availability | Triggered when a capacity disk used for vSAN is smaller than 255 GB (default max component size), so virtual machines that run on the vSAN Datastore might experience disk space issues. |
| Controller with pass-through and RAID disks has issues. | Storage | Configuration | Triggered when a controller with pass-through and RAID disks has issues. |
| Disk format version of one or more vSAN disks is out of date | Storage | Configuration | Triggered when the disk format version of one or more vSAN disks is out of date and is not compatible with other vSAN disks. This can lead to problems in creating or powering on VMs, performance degradation, and EMM failures. |
| ESXi host issues retrieving hardware info. | Storage | Configuration | Triggered when the ESXi host issues retrieving hardware info. |
| Firmware provider hasn't all its dependencies met or is not functioning as expected. | Storage | Configuration | Triggered when a firmware provider has not met all its dependencies or is not functioning as expected. |
| Host with inconsistent extended configurations is detected. | Storage | Configuration | Triggered when a host with inconsistent extended configurations is detected. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| | | | vSAN cluster extended configurations are set as object repair timer is 60 minutes, site read locality is Enabled, customized swap object is Enabled, large scale cluster support is Disabled; For host with inconsistent extended configurations, vSAN cluster remediation is recommended, for host doesn't support any extended configuration, ESXi software upgrade is needed; And to make cluster scalability configuration take effect, host reboot could be required. |
| Inconsistent configuration (like dedup/compression, encryption) setup on hosts or disks with the cluster. | Storage | Configuration | Triggered when there is inconsistent configuration (like dedup/compression, encryption) setup on hosts or disks with the cluster. |
| Network adapter driver is not VMware certified. | Storage | Configuration | Triggered when the network adapter driver is not VMware certified. |
| Network adapter firmware is not VMware certified. | Storage | Configuration | Triggered when the network adapter firmware is not VMware certified. |
| Network adapter is not VMware certified. | Storage | Configuration | Triggered when the network adapter is not VMware certified. |
| Network configuration of the vSAN iSCSI target service is not valid. | Storage | Availability | Triggered when the network configuration of the vSAN iSCSI target service is not valid.<br><br>This health check validates the presence of the default vmknic for the vSAN iSCSI target service, and verifies that all the existing targets have valid vmknic configurations. |
| Non-vSAN disks are used for VMFS or Raw Device Mappings(RDMs). | Storage | Availability | Triggered when non-vSAN disks are used for VMFS or Raw Device Mappings (RDMs). |
| Number of vSAN components on a disk is reaching or has reached its limit. | Storage | Capacity | Triggered when the number of vSAN components on a disk is reaching or has reached its limit. This will cause failure in the deployment of new Virtual Machines and also impact rebuild operations. |
| Number of vSAN components on a host is reaching or has reached its limit. | Storage | Capacity | Triggered when the number of vSAN components on a host is reaching or has reached its limit.<br><br>This will cause failure in the deployment of new Virtual Machines and also impact rebuild operations. |
| One or more ESXi hosts in the cluster do not support CPU AES-NI or have it disabled. | Storage | Availability | Triggered when one or more hosts in the cluster do not support CPU AES-NI or have it disabled. As a result, the system might use the software encryption that is significantly slower than AES-NI. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| RAID controller configuration has issues. | Storage | Configuration | Triggered when the RAID controller configuration has issues. |
| Storage I/O controller driver is not VMware certified | Storage | Configuration | Triggered when the stability and integrity of vSAN may be at risk as the storage I/O controller driver is not VMware certified. |
| Storage I/O controller drivers is not supported with the current version of ESXi running on the host | Storage | Configuration | Triggered when the stability and integrity of vSAN may be at risk as the storage I/O controller driver is not supported with the current version of ESXi running on the host. |
| Storage I/O Controller firmware not is VMware certified. | Storage | Configuration | Triggered when the storage I/O Controller firmware not is VMware certified. |
| Storage I/O controller is not compatible with the VMware Compatibility Guide | Storage | Configuration | Triggered when the vSAN environment may be at risk as the Storage I/O controller on the ESXi hosts that are participating in a vSAN cluster are not compatible with the VMware Compatibility Guide. |
| The current status of the Customer Experience Improvement Program (CEIP) not is enabled. | Storage | Availability | Triggered when the current status of the Customer Experience Improvement Program (CEIP) not is enabled. |
| The Internet connectivity is not available for vCenter Server. | Storage | Availability | Triggered when internet connectivity is not available for vCenter Server. |
| The resync operations are throttled on any hosts. | Storage | Configuration | Triggered when resync operations are throttled. Please clear the limit, unless you need it for particular cases like a potential cluster meltdown. |
| Time of hosts and VC are not synchronized within 1 minute. | Storage | Configuration | Triggered when the time of hosts and VC are not synchronized within 1 minute.<br><br>Any difference larger than 60 seconds will lead this check to fail. If the check fails, it is recommended that you check the NTP server configuration. |
| vCenter Server or any of the ESXi hosts experience problems when connecting to Key Management Servers (KMS). | Storage | Availability | Triggered when the vCenter Server or any of the hosts experience problems when connecting to KMS. |
| vCenter server state was not pushed to ESXi due to vCenter server being out of sync. | Storage | Configuration | Triggered when the vCenter server state was not pushed to ESXi due to vCenter server being out of sync.<br><br>During normal operation, the vCenter server state is regarded as source of truth, and ESXi hosts are automatically updated with the latest host membership list. When vCenter server is replaced or recovered from backup, the host membership list in vCenter server may be out of sync. This health check detects such cases, and alerts if vCenter server state was not |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| | | | pushed to ESXi due to vCenter server being out of sync. In such cases, first fully restore the membership list in vCenter server, and then perform 'Update ESXi configuration' action if required. |
| vSAN and VMFS datastores are on a same Dell H730 controller with the lsi_mr3driver. | Storage | Configuration | Triggered when the vSAN and VMFS datastores are on a same Dell H730 controller with the lsi_mr3driver. |
| vSAN build recommendation based on the available releases and VCG compatibility guide. | Storage | Availability | Triggered when the vSAN build is not compatible with available releases and VCG compatibility guide.<br><br>This is the ESXi build that vSAN recommends as the most appropriate, given the hardware, its compatibility per the VMware Compatibility Guide and the available releases from VMware. |
| vSAN build recommendation engine has all its dependencies met and is functioning as expected. | Storage | Availability | Triggered when the vSAN build recommendation engine has issues.<br><br>The vSAN Build Recommendation Engine relies on the VMware compatibility guide and VMware release metadata for its recommendation. To provide build recommendations, it also requires VMware Update Manager service availability, internet connectivity, and valid credentials for my.vmware.com. This health check ensures that all dependencies are met and that the recommendation engine is functioning correctly. |
| Some disk(s) free space in vSAN Cluster is less than 10% | Storage | Capacity | Triggered when the disk usage in a vSAN cluster reaches 90% of capacity.<br><br>Cleared by removing virtual machines that are no longer in use or adding more disks to the cluster. |
| Some disk(s) free space in vSAN Cluster is less than 30% | Storage | Capacity | Triggered when the disk usage in a vSAN cluster reaches 70% of capacity.<br><br>Cleared by removing virtual machines that are no longer in use or adding more disks to the cluster. |
| vSAN cluster is reaching or has reached its limit for components, free disk space and read cache reservations. | Storage | Capacity | Triggered when the vSAN cluster is reaching or has reached its limit for components, free disk space and read cache reservations. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN Cluster virtual disk count capacity is less than 5%. | Storage | Capacity | Triggered when the number of virtual disks per host in the vSAN cluster reaches 95% of capacity.<br><br>Cleared by adding most hosts to the cluster. |
| vSAN Cluster virtual disk count is approaching capacity. | Storage | Capacity | Triggered when the number of virtual disks per host in the vSAN cluster reaches 75% of capacity.<br><br>Cleared by adding most hosts to the cluster. |
| vSAN configuration for LSI 3108-based controller has issues. | Storage | Configuration | Triggered when the vSAN configuration for LSI 3108-based controller has issues. |
| vSAN disk group type (All-Flash or Hybrid) for the used SCSI controller is not VMware certified. | Storage | Configuration | Triggered when the vSAN disk group type (All-Flash or Hybrid) for the used SCSI controller is not VMware certified. |
| vSAN enabled hosts have inconsistent values for advanced configuration options. | Storage | Configuration | Triggered when some advanced configuration settings have different values on different hosts in the vSAN cluster. |
| vSAN firmware version recommendation based on the VCG. | Storage | Configuration | Triggered when the vSAN firmware version recommendation based on the VCG check has issues. |
| vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk. | Storage | Availability | Triggered when the vSAN has encountered an integrity issue with the metadata of an individual component on a physical disk. |
| vSAN HCL DB auto updater is not working properly. | Storage | Configuration | Triggered when the vSAN HCL DB auto updater is not working properly. This means that vSAN cannot download and update its HCL DB automatically. |
| vSAN HCL DB is not up-to-date. | Storage | Configuration | Triggered when the vSAN HCL DB is not up-to-date. |
| vSAN Health Service is not able to find the appropriate controller utility for the storage controller on the ESXi host. | Storage | Availability | Triggered when the vSAN Health Service is not able to find the appropriate controller utility for the storage controller on the ESXi host. |
| vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks. | Storage | Performance | Triggered when the vSAN is running low on the vital memory pool (heaps) needed for the operation of physical disks.<br><br>This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive. |
| vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks. | Storage | Performance | Triggered when the vSAN is running low on the vital memory pool (slabs) needed for the operation of physical disks. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| | | | This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive. |
| vSAN is using a physical disk which has high congestion value. | Storage | Performance | Triggered when the vSAN is using a physical disk which has high congestion value.<br><br>This can lead to a variety of performance issues such as virtual machine storage performance degradation, operation failures, or even ESXi hosts going unresponsive. |
| vSAN iSCSI target service home object has issues. | Storage | Availability | Triggered when the vSAN iSCSI target service home object has issues.<br><br>This health check verifies the integrity of the vSAN iSCSI target service home object. It also verifies that the configuration of the home object is valid. |
| vSAN iSCSI target service is not running properly or is not correctly enabled on the host. | Storage | Availability | Triggered when the vSAN iSCSI target service is not running properly or is not correctly enabled on the host.<br><br>This health check verifies the service runtime status of the vSAN iSCSI target service, and checks whether the service is correctly enabled on each host. |
| vSAN performance service statistics database object is reporting issues. | Storage | Availability | Triggered when the vSAN performance service statistics database object is reporting issues. |
| vSphere cluster members do not match vSAN cluster members. | Storage | Configuration | Triggered when the vSphere cluster members do not match vSAN cluster members. |

**Table 523: vSAN Cluster Object Efficiency Alert Definitions**

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN Cluster flash read cache is approaching capacity. | Storage | Capacity | Triggered when the Read Cache (RC) in the vSAN cluster reaches 80% of capacity.<br><br>Cleared by adding flash storage to the read cache. |
| vSAN Cluster flash read cache capacity is less than 5%. | Storage | Capacity | Triggered when the Read Cache (RC) in the vSAN cluster reaches 95% of capacity.<br><br>Cleared by adding flash storage to the read cache. |

## vSAN Adapter Instance Object Alert Definitions

Alerts on the vSAN Adapter Instance Object have health impact.

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues. | Storage | Configuration | Triggered when the vSAN adapter instance failed to collect data from vSAN Health Service. The health Service might have issues. |

**vSAN Disk Group Object Alert Definitions**

Alerts on the vSAN Disk Group Object have efficiency impact.

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN Disk Group read cache hit rate is less than 90%. | Storage | Performance | Triggered when the vSAN disk group read cache hit rate is less than 90%.<br><br>Cleared by adding more cache to accommodate the workload. |
| vSAN Disk Group read cache hit rate is less than 90% and write buffer free space is less than 10%. | Storage | Capacity | Triggered when the vSAN disk group read cache hit rate is less than 90% and the vSAN disk group write buffer free space is less than 10%.<br><br>Cleared by adding more flash capacity to the vSAN disk group. |

**vSAN Host Object Alert Definitions**

Alerts on the vSAN Host Object have security impact.

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN host has encryption disabled, while the vSAN cluster has encryption enabled. | Storage | Configuration | Triggered when the vSAN host has encryption disabled, while the vSAN cluster has encryption enabled.<br><br>Cleared by enabling encryption on vSAN host. |
| vSAN host encryption is enabled, while the vSAN cluster encryption is disabled. | Storage | Configuration | Triggered when the vSAN host has encryption enabled, while the vSAN cluster has encryption disabled.<br><br>Cleared by enabling encryption on vSAN cluster. |

**vSAN Capacity Disk Object Alert Definitions**

Alerts on the vSAN Capacity Disk object have security impact.

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN capacity disk has encryption disabled, while the vSAN cluster has encryption enabled. | Storage | Configuration | Triggered when the vSAN capacity disk has encryption disabled, while the vSAN cluster has encryption enabled.<br><br>Cleared by enabling encryption on vSAN capacity disk. |
| vSAN capacity disk encryption is enabled, while the vSAN cluster encryption is disabled. | Storage | Configuration | Triggered when the vSAN capacity disk has encryption enabled, while the vSAN cluster has encryption disabled.<br><br>Cleared by enabling encryption on vSAN cluster. |
| The free read cache reservations across the entire vSAN cluster are beyond the thresholds. | Storage | Capacity | Triggered when the flash read cache is exhausted.<br><br>**NOTE**<br>Flash read cache is only relevant to hybrid configurations and is not relevant on all-flash configurations. |
| Deployment of new virtual machines fails due to insufficient disk capacity | Storage | Capacity | Triggered when the disk capacity of the vSAN cluster exceeds the threshold value. |

**vSAN Cache Disk Object Alert Definitions**

Alerts on the vSAN Cache Disk object have security impact.

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN cache disk has encryption disabled, while the vSAN cluster has encryption enabled. | Storage | Configuration | Triggered when the vSAN cache disk has encryption disabled, while the vSAN cluster has encryption enabled.<br><br>Cleared by enabling encryption on vSAN cache disk. |
| vSAN cache disk encryption is enabled, while the vSAN cluster encryption is disabled. | Storage | Configuration | Triggered when the vSAN cache disk has encryption enabled, while the vSAN cluster has encryption disabled.<br><br>Cleared by enabling encryption on vSAN cluster. |

**vSAN File Service Alert Definitions**

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN File Service infrastructure health has issues. | Storage | Configuration | Triggered when there is an issue with file service infrastructure health state of an ESXi host in the vSAN cluster. |

*Table continued on next page*

*Continued from previous page*

| Alert | Alert Type | Alert Subtype | Description |
|---|---|---|---|
| vSAN File Share health is not in a good state. | Storage | Configuration | Triggered when the vSAN File Share health is not in a good state. |
| Network File System (NFS) daemon is not running. | Storage | Configuration | Triggered when the NFS daemon process is not running. |
| Root File System is inaccesible. | Storage | Configuration | Triggered when the root file system does not repond to the file server. |
| File Server IP address not assigned. | Storage | Configuration | Triggered when IP address is not assigned to the file server. |
| vSAN File Server health is not in a good state. | Storage | Configuration | Triggered when the vSAN File Server health is not in a good state. |

### Alerts in the vSphere Web Client

The vSphere Web Client displays the results of health tests for the following vSAN monitored groups:

- Network
- Physical disk
- Cluster
- Limits
- Data
- Hardware compatibility
- Performance Service
- Stretched Cluster (if enabled)

Each group contains several individual checks. If a check fails, the vSAN adapter issues a warning or error level alert. The alert indicates the host or cluster where the problem occurred and provides a recommendation to clear the alert. For a complete list of all vSAN health test alerts, see Knowledge Base article 2114803.

### vSphere Distributed Port Group

The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Port objects in your environment.

#### Health/Critical

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Critical

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| One or more ports are in a link down state. | Symptoms include all of the following:<br>• Port is connected. | Verify that there is physical connectivity for the NICs on the host. Verify the admin status on the port. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | • One or more ports are in a link down state. | |
| One or more ports are experiencing network contention. | Port is experiencing dropped packets. | Check if the packet drops are due to high CPU resource utilization or uplink bandwidth utilization. User vMotion to migrate the virtual machine that the port is attached to a different host. |

**Virtual Machine Alert Definitions**

The vCenter adapter provides alert definitions that generate alerts on the virtual machine objects in your environment.

**Health/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Symptom-based

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Virtual machine is experiencing memory compression, ballooning or swapping due to memory limit. | • Virtual machine memory limit is set AND<br>• Virtual machine memory demand exceeds configured memory limit AND<br>• [Virtual machine memory is compressed OR<br>• Virtual machine is using swap OR<br>• Virtual machine memory ballooning is at warning/immediate/critical level] AND<br>• Recommended virtual machine memory size | Increase the memory limit for the virtual machine to match the recommended memory size. Alternatively, remove memory limit for the virtual machine. |
| Virtual machine has CPU contention caused by IO wait. | Virtual machine CPU I/O wait is at warning/immediate/critical level. | Increase the datastore I/O capacity for the connected data stores to reduce CPU I/O wait on the virtual machine. |
| Virtual machine has unexpected high memory workload. | Symptoms include all of the following:<br>• Virtual machine memory workload is at Warning/Immediate/Critical level<br>• Anomaly is starting to/moderately/ critically high | 1. Check the guest applications to determine whether high memory workload is an expected behavior.<br>2. Add more memory for this virtual machine. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Virtual machine has memory contention due to swap wait and high disk read latency. | Symptoms include all of the following:<br>• Virtual machine CPU swap wait is at warning/immediate/critical level (5/10/15)<br>• Virtual machine has read latency at warning level<br>• Recommended virtual machine memory size | Add more memory for this virtual machine. |
| Virtual machine has memory contention due to memory compression, ballooning or swapping. | • ! Virtual machine memory limit is set AND<br>• Virtual machine has memory contention at warning/immediate/ critical level AN<br>• [ Virtual machine memory ballooning at warning/immediate/ critical level OR<br>• Virtual machine memory is compressed OR<br>• Virtual machine is using swap] | 1. Add memory reservations to this virtual machine to prevent ballooning and swapping.<br>2. Use vSphere vMotion to migrate this virtual machine to a different host or cluster. |
| Virtual machine has disk I/O read latency problem. | Symptoms include all of the following:<br>• Virtual machine disk read latency at Warning /Immediate/Critical level<br>• Virtual machine disk read latency above DT<br>• Virtual machine has low co-stop<br>• Virtual machine has low CPU swap wait | 1. Check whether you have enabled Storage IO control on the datastores connected to the virtual machine.<br>2. Increase IOPS for the datastores connected to the virtual machine.<br>3. Use vSphere Storage vMotion to migrate this virtual machine to a different datastore with higher IOPS. |
| Virtual machine has disk I/O write latency problem. | Symptoms include all of the following:<br>• Virtual machine disk write latency at Warning/ Immediate/Critical level<br>• Virtual machine disk write latency above DT<br>• Virtual machine has low CPU swap wait (< 3 ms) | 1. Check whether you have enabled Storage IO Control on the data stores connected to the datastore.<br>2. Increase IOPS for the data stores connected to the virtual machine.<br>3. If the virtual machine has multiple snapshots, delete the older snapshots.<br>4. Use vSphere Storage vMotion to migrate some virtual machines to a different datastore. |
| Virtual machine has disk I/O latency problem caused by snapshots. | Symptoms include all of the following:<br>• Virtual machine CPU I/O wait is at warning/immediate/critical level<br>• Virtual machine has at least one snapshot<br>• All child datastores have [ ! Disk command latency at warning level ] | 1. If the virtual machine has multiple snapshots, delete the older snapshots.<br>2. Reduce the number of snapshots by consolidating the snapshots into one snapshot. In vSphere Client, |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | | select the VM, right-click, select **Snapshot**, and then **Consolidate**. |
| Not enough resources for vSphere HA to start the virtual machine. | Not enough resources for vSphere HA to start VM (Fault symptom). | 1. If virtual machine CPU reservation is set, decrease the CPU reservation configuration.<br>2. If virtual machine memory reservation is set, decrease the memory reservation configuration.<br>3. Add more hosts to cluster.<br>4. Bring any failed hosts online or resolve a network partition, if one exists.<br>5. If DRS is in manual mode, look for pending recommendations and approve the recommendations so that vSphere HA failover can proceed. |
| The Fault tolerance state of the virtual machine has changed to "Disabled" state. | VM fault tolerance state changed to disabled (Fault symptom). | Enable the secondary virtual machine indicated in the alert. |
| vSphere HA failed to restart a network isolated virtual machine. | vSphere HA failed to restart a network isolated virtual machine (Fault symptom). | Manually power on the virtual machine. |
| The fault tolerance state of the virtual machine has changed to "Needs Secondary" state. | VM Fault Tolerance state changed to needs secondary (Fault symptom). | Keep HA enabled when Fault tolerance (FT) is required to protect virtual machines. |
| vSphere HA cannot perform a failover operation for the virtual machine | vSphere HA virtual machine failover unsuccessful (Fault symptom) | 1. If the error information reports that a file is locked, the virtual machine might be powered on a host that the vSphere HAprimary agent can no longer monitor by using the management network or heartbeat datastores.<br>2. The virtual machine might have been powered on by a user on a host outside of the cluster. If any hosts are declared offline, determine whether a networking or storage problem caused the situation.<br>3. If the error information reports that the virtual machine is in an invalid state, an in-progress operation might be preventing access to the virtual machine files. Determine whether any operations are in progress, such as a clone |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | | operation that is taking a long time to complete.<br>4. You can also try to power on the virtual machine and investigate any returned errors. |
| One or more virtual machine guest file systems are running out of disk space. | Symptoms include all of the following:<br>• Guest file system usage at warning level<br>• Guest file system usage at critical level | Add a new virtual hard disk or expand the existing disk of the virtual machine. Before expanding the existing disk, remove all the snapshots. Once done, use a guest OS specific procedure to expand the file system on the new or expanded disk. |
| Virtual machine has CPU contention due to memory page swapping in the host. | Symptoms include all of the following:<br>• Virtual machine CPU swap wait is at Critical level<br>• Virtual machine CPU swap wait is at Immediate level<br>• Virtual machine CPU swap wait is at Warning level | 1. Set memory reservations for the virtual machine to prevent its memory from being swapped.<br>2. Verify that VMware Tools is installed and running, and that the balloon driver is enabled in the guest. Memory ballooning helps the host reclaim unused memory from the guest more effectively, and might avoid swapping.<br>3. Use vMotion to migrate this virtual machine to a different host or cluster. |

**Efficiency/Warning**

These alert definitions have the following impact and criticality information.

**Impact**

Efficiency

**Criticality**

Warning

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Virtual machine is idle. | Symptoms include all of the following:<br>• Virtual machine is idle<br>• Virtual machine high ready time on each vCPU<br>• ! Virtual machine is powered off | Power off this virtual machine to allow for other virtual machines to use CPU and memory that this virtual machine is wasting. |

**Risk/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Risk

**Criticality**

Symptom-based

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Virtual machine has CPU contention caused by co-stop. | Symptoms include all of the following:<br>• Virtual machine CPU co-stop at warning/immediate/critical level<br>• ! Virtual machine is powered off<br>• Number of vCPUs to remove from virtual machine | Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended by the symptom. |
| Virtual machine is violating vSphere 5.5 hardening guide. | • Unrestricted VM-to-VM communication through VMCI OR<br>• VMsafe CPU/Memory APIs-port number configured OR<br>• Dvfilter network API enabled OR<br>• Non-compliant max VMX file size OR<br>• Non-compliant max VM log file size OR<br>• Allow unauthorized modification of device settings OR<br>• Allow unauthorized connect and disconnect of devices OR<br>• Tools auto install not disabled OR<br>• Non-compliant max number of remote console connections OR<br>• Allow VM to obtain detailed information about the physical host OR<br>• Non-compliant max VM log file count OR<br>• Feature not exposed in vSphere: MemsFss is not disabled OR<br>• VMsafe CPU/memory API enabled OR<br>• Parallel port connected OR<br>• Console drag and drop operation not disabled OR<br>• Console copy operation not disabled OR<br>• Serial port connected OR<br>• Feature not exposed in vSPhere: AutoLogon is not disabled OR<br>• Use independent non persistent disk OR | Fix the vSphere 5.5 hardening guide rule violations according to the recommendations in the vSphere Hardening Guide (XLSX). |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | • Feature not exposed in vSphere: UnityPush is not disabled OR<br>• Shrink virtual disk not disabled - diskShrink OR<br>• Feature not exposed in vSphere: GetCreds is not disabled OR<br>• CD-ROM connected OR<br>• Feature not exposed in vSPhere: HGFSServerSet is not disabled OR<br>• Console paste operation not disabled OR<br>• Feature not exposed in vSPhere: BIOSBBS is not disabled OR<br>• Shrink virtual disk not disabled - diskWiper OR<br>• USB controller connected OR<br>• Feature not exposed in vSPhere: Monitor Control is not disabled OR<br>• Floppy drive connected OR<br>• Feature not exposed in vSphere: LaunchMenu is not disabled OR<br>• Versionget is not disabled OR<br>• Feature not exposed in vSphere: Toporequest is not disabled OR<br>• Feature not exposed in vSphere: Unity-interlock not disabled OR<br>• VM logging is not disabled OR<br>• Feature not exposed in vSphere: Unity is not disabled OR<br>• Feature not exposed in vSphere: Trashfolderstate is not disabled OR<br>• VGA only mode is not enabled OR<br>• Feature not exposed in vSphere: Trayicon is not disabled OR<br>• Feature not exposed in vSphere: Unity-Taskbar is not disabled OR<br>• Feature not exposed in vSphere: Versionset is not disabled OR<br>• VM console access via VNC protocol is not disabled OR<br>• Feature not exposed in vSphere: Protocolhandler is not disabled OR<br>• VIX message is not disabled OR<br>• Feature not exposed in vSphere: Shellaction is not disabled OR<br>• 3D features is not disabled OR | |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | • Feature not exposed in vSphere: Unity-Windowcontents is not disabled OR<br>• Feature not exposed in vSphere: Unity-Unityactive is not disabled | |
| Virtual machine has CPU contention due to multi-vCPU scheduling issues (co-stop) caused by snapshots | Symptoms include all of the following:<br>• Virtual machine CPU co-stop is at Warning level OR<br>• Virtual machine CPU co-stop is at Immediate level OR<br>• Virtual machine CPU co-stop is at Critical level<br>And<br>• Virtual machine is powered off OR<br>• Virtual machine has at least one snapshot | None. |

### vSphere Distributed Switch Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vSphere Distributed Switch objects in your environment.

#### Health/Critical

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Critical

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Network traffic is blocked for one or more ports. | Network traffic is blocked for one or more ports. | Check the security policy on the port groups as well as any ACL rule configuration. |

#### Health/Warning

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Warning

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Distributed Switch configuration is out of sync. | Distributed Switch configuration is out of sync with the vCenter. | Change the distributed switch configuration to match the host. Identify the distributed switch properties that are out of sync. If these properties were changed locally on the host in order to maintain connectivity, update the distributed switch configuration in the vCenter. Otherwise, re-apply the the vCenter configuration to this host. |
| One or more VLANs are unsupported by the physical switch. | One or more VLANs are unsupported by the physical switch. | Ensure the VLAN configuration on the physical switch and the distributed port groups are consistent. |
| Teaming configuration does not match the physical switch. | Teaming configuration does not match the physical switch. | Ensure the teaming configuration on the physical switch and the distributed switch are consistent. |
| The MTU on the Distributed Switch is not allowed by one or more VLANs on the host. | The MTU on the Distributed Switch is not allowed by one or more VLANs on the host. | Ensure the MTU configuration on the physical switch and the distributed switch are consistent. |
| There is an MTU mismatch between the host and a physical switch. | There is an MTU mismatch between the host and a physical switch. | Adjust the MTU configuration on the host to match the physical switch. Change the MTU configuration on the physical switch. |

**Risk/Warning**

These alert definitions have the following impact and criticality information.

**Impact**

Risk

**Criticality**

Warning

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| The distributed switch configuration is incorrect. | Host without redundant physical connectivity to the distributed switch. | Verify that at least two NICs on each host is connected to the distributed switch. |

### vCenter Server Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vCenter objects in your environment.

**Health/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Symptom-based

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| A problem occurred with a vCenter component. | The vCenter health changed (fault symptom). | The actions to take to resolve the problems depend on the specific problem that caused the fault. Review the issue details, and check the documentation. |
| Duplicate object name found in the vCenter. | Duplicate object name found in the vCenter. | Ensure that the virtual machines names are unique before enabling the Name-Based Identification feature. |
| The vCenter Storage data collection failed. | The vCenter storage data collection failed. | Ensure vCenter Management Webservice is started and Storage Management Service is functioning. |
| VASA Provider(s) disconnected | One or more VASA Providers disconnected from vCenter. | If the VASA provider is inaccessible from the vCenter and you are getting an invalid certificate error then, see KB article:2079087. Contact the hardware vendor for further support. |
| Certificate for VASA Provider(s) will expire soon | One or more VASA Providers' certificates expire soon. | Contact the hardware vendor for getting support on the CA certificates and CRLs for VASA provider. |
| Refreshing CA certificates and CRLs for VASA Provider(s) failed | Refreshing CA certificates and CRLs for one or more VASA Providers failed. | Refresh the storage provider certificate as per the following document: *Refresh Storage Provider Certificates*. Contact the hardware vendor for further support.<br>**NOTE**<br>The *Refresh Storage Provider Certificates* is in the vSphere Storage 6.5 guide. |
| Virtual machine has memory contention caused by swap wait and high disk read latency. | Virtual Machine has a memory contention due to swap wait and high disk read latency. | Add more memory for the virtual machine and ensure that VMware Tools is running in the virtual machine. |

**Risk/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Risk

**Criticality**

Symptom-based

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Virtual machine has CPU contention due to multi-vCPU scheduling issues (co-stop) caused by too many vCPUs. | Virtual Machine experiences a high co-stop. The co-stop is the amount of time taken when the virtual machine is ready to run but is experiencing delay because of the co-vCPU scheduling contention. High co-stop occurs when too many vCPUs are configured for the virtual machine, and not enough physical CPUs are available to manage the co-vCPU scheduling. | Review the symptoms listed and remove the number of vCPUs from the virtual machine as recommended. |

### Datastore Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the datastore objects in your environment.

#### Health/Critical

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Critical

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| A storage device for a datastore has been detected to be off. | Storage device has been turned off administratively (fault symptom). | Ask the administrator about the device state. The fault will be resolved and the alert canceled if the device is turned on. If SCSI devices are detached or permanently removed, you must manually cancel the alert. |
| Datastore has lost connectivity to a storage device. | Host(s) lost connectivity to storage device(s) (fault symptom). | The storage device path, for example, `vmhba35:C1:T0:L7`, contains several potential failure points: Path Element \| Failure Point<br>---<br>vmhba35 \| HBA (Host Bus Adapter) C1 \| Channel T0 \| Target (storage processor port) L7 \| LUN (Logical Unit Number or Disk Unit).<br>To determine the cause of the failure or to eliminate possible problems: Identify the available storage paths to the reported storage device by running **`esxcfg-mpath -l`**. For more information, see http://kb.vmware.com/kb/1003973. Check that a rescan does not restore visibility to the targets. For information on rescanning the storage |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | | device by using the command-line interface and the vSphere Client, see http://kb.vmware.com/kb/1003988. Determine whether the connectivity issue is with the iSCSI storage or the fiber storage.<br><br>Troubleshoot the connectivity to the iSCSI storage by using the software initiator:<br>1. Check whether a ping to the storage array fails from ESX. For more information, see http://kb.vmware.com/kb/1003486<br>2. Check whether a vmkping to each network portal of the storage array fails. For more information, seehttp://kb.vmware.com/kb/10037828.<br>3. Check that the initiator is registered on the array. For more information, contact your storage vendor.<br>4. Check that the following physical hardware is functioning correctly: Ethernet switch, Ethernet cables between the switch and the ESX host, and Ethernet cables between the switch and the storage array.<br>To troubleshoot the connectivity to the fiber-attached storage, check the fiber switch. The fiber switch zoning configuration permits the ESX host to see the storage array. If you require assistance, contact your switch vendor. The fiber switch propagates RSCN messages to the ESX hosts. For more information about configuring the fiber switch, see http://kb.vmware.com/kb/1002301.<br><br>Finally, check the following physical hardware: the storage processors on the array, the fiber switch and the Gigabit Interface Converter (GBIC) units in the switch, the fiber cables between the fiber switch and the array, and the array itself.<br><br>You must rescan after making changes to make sure that the targets are |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | | detected. If storage connectivity is restored for all of the affected host and storage device combinations, the fault is cleared and the alert canceled. If storage connectivity for the devices indicated is caused by a permanent loss or change, you must cancel the fault alert as a workaround. The alert will then be canceled automatically. |

**Health/Immediate**

These alert definitions have the following impact and criticality information.

**Impact**

Health

**Criticality**

Immediate

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Datastore has one or more hosts that have lost redundant paths to a storage device. | Host(s) lost redundancy to storage device(s) (fault symptom). | The storage device path, for example, vmhba35:C1:T0:L7, contains several potential failure points:<br>Path Element \| Failure Point<br>---<br>vmhba35 \| HBA (Host Bus Adapter) C1 \| Channel T0 \| Target (storage processor port) L7 \| LUN (Logical Unit Number or Disk Unit).<br><br>Use the following guidance to determine the cause of the failure or to eliminate possible problems. Identify the available storage paths to the reported storage device by running **`esxcfg-mpath -l`**. For more information, see http://kb.vmware.com/kb/1003973.<br><br>Check that a rescan does not restore visibility to the targets. For information on rescanning the storage device by using the command-line interface and the vSphere Client, see http://kb.vmware.com/kb/1003988. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | | Determine whether the connectivity issue is with the iSCSI storage or the fiber storage. Troubleshoot the connectivity to the iSCSI storage by using the software initiator:<br>1. Check whether a ping to the storage array fails from ESX. For more information, seehttp://kb.vmware.com/kb/1003486.<br>2. Check whether a vmkping to each network portal of the storage array fails. For more information, seehttp://kb.vmware.com/kb/10037828.<br>3. Check that the initiator is registered on the array. For more information, contact your storage vendor.<br>4. Check that the following physical hardware is functioning correctly: Ethernet switch, Ethernet cables between the switch and the ESX host, and Ethernet cables between the switch and the storage array.<br>To troubleshoot the connectivity to the fiber-attached storage, check the fiber switch. The fiber switch zoning configuration permits the ESX host to see the storage array. If you require assistance, contact your switch vendor. The fiber switch propagates RSCN messages to the ESX hosts. For more information about configuring the fiber switch, seehttp://kb.vmware.com/kb/1002301.<br><br>Finally, check the following physical hardware: the storage processors on the array, the fiber switch and the Gigabit Interface Converter (GBIC) units in the switch, the fiber cables between the fiber switch and the array, and the array itself. You must rescan after making changes to make sure that the targets are detected. If storage connectivity is restored for all of the affected host and storage device combinations, the fault is cleared and the alert canceled. If storage connectivity for the devices indicated is caused by a permanent loss or |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| | | change, you must cancel the fault alert as a workaround. The alert will be canceled automatically after that. |

**Risk/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Risk

**Criticality**

Symptom-based

| Alert Definition | Symptom | Recommendations |
|---|---|---|
| Datastore is running out of disk space. | Symptoms include all of the following:<br>• Datastore space usage reaching warning/immediate/critical level<br>• ! Datastore space growth above DT<br>• Datastore space time remaining is low | 1. Add more capacity to the datastore.<br>2. Use vSphere vMotion to migrate some virtual machines to a different datastore.<br>3. Delete unused snapshots of virtual machines from datastore.<br>4. Delete any unused templates on the datastore. |

**Data Center Alert Definitions**

The vCenter adapter provides alert definitions that generate alerts on the Data Center objects in your environment.

**Risk/Symptom-Based**

These alert definitions have the following impact and criticality information:

**Impact**

Risk

**Criticality**

Symptom-based

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Data center has unbalanced CPU "demand" workload. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• DC is unbalanced on CPU "demand" workload | Rebalance the container to spread the workload more evenly. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | • DC has significant CPU "demand" workload difference<br>• At least one cluster in DC has high CPU "demand" workload | |
| Data center has unbalanced memory "demand" workload. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully enabled<br>• DC is unbalanced on memory "demand" workload difference<br>• At least one cluster in DC has high memory "demand" workload | Rebalance the container to spread the workload more evenly. |
| Data center has unbalanced memory "consumed" workload. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• DC is unbalanced on memory "consumed" workload<br>• DC has significant memory "consumed" workload difference<br>• At least one cluster in DC has high memory "consumed" workload | Rebalance the container to spread the workload more evenly. |

**Custom Data Center Alert Definitions**

The vCenter adapter provides alert definitions that generate alerts on the Custom Data Center objects in your environment.

**Risk/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Risk

**Criticality**

Symptom-based

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Custom data center has unbalanced CPU "demand" workload. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• CDC is unbalanced on CPU "demand" workload<br>• CDC has significant CPU "demand" workload difference<br>• At least one cluster in CDC has high CPU "demand" workload | Rebalance the container to spread the workload more evenly. |
| Custom data center has unbalanced memory "demand" workload. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• CDC is unbalanced on memory "demand" workload<br>• CDC has significant memory "demand" workload difference<br>• At least one cluster in CDC has high memory"demand" workload | Rebalance the container to spread the workload more evenly. |
| Custom Datacenter has unbalanced memory "consumed" workload. | Symptoms include all of the following:<br>• DRS enabled<br>• DRS fully automated<br>• CDC is unbalanced on memory "consumed" workload<br>• CDC has significant memory "consumed" workload difference<br>• At least one cluster in CDC has high memory"consumed" workload | Rebalance the container to spread the workload more evenly. |

### vSphere Pod Alert Definitions

The vCenter adapter provides alert definitions that generate alerts on the vSphere Pod objects in your environment.

**Health/Symptom-Based**

These alert definitions have the following impact and criticality information.

**Impact**

Risk/Health

**Criticality**

Symptom-based

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Not enough resources for vSphere HA to start the Pod | Not enough resources for vSphere HA to start Pod | |
| One or more Pod guest file systems are running out of disk space | Symptom set is met when any of the symptoms are true: | |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | • Guest file system space usage at warning level<br>• Guest file system space usage at critical level | |
| Pod CPU usage is at 100% for an extended period of time | Pod sustained CPU usage is 100% | |
| Pod disk I/O read latency is high | Symptom set is met when any of the symptoms are true:<br>• Pod disk read latency at Warning level<br>• Pod disk read latency at Immediate level<br>• Pod disk read latency at Critical level | |
| Pod disk I/O write latency is high | Symptom set is met when any of the symptoms are true:<br>• Pod disk write latency at Warning level<br>• Pod disk write latency at Immediate level<br>• Pod disk write latency at Critical level | |
| Pod has CPU contention due to long wait for I/O events | Symptom set is met when any of the symptoms are true:<br>• Pod CPU I/O wait is at Critical level<br>• Pod CPU I/O wait is at Immediate level<br>• Pod CPU I/O wait is at Warning level | |
| Pod has CPU contention due to memory page swapping in the host | Symptom set is met when any of the symptoms are true.<br>• Pod CPU swap wait is at Critical level<br>• Pod CPU swap wait is at Immediate level<br>• Pod CPU swap wait is at Warning level | |
| Pod has CPU contention due to multi-vCPU scheduling issues (co-stop) caused by too many vCPUs | Alert is triggered when all of the symptom sets are true.<br>• Pod is powered off<br><br>Symptom set is met when any of the symptoms are true.<br>• Pod CPU co-stop is at Critical level<br>• Pod CPU co-stop is at Immediate level | |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | • Pod CPU co-stop is at Warning level | |
| Pod has memory contention caused by swap wait and high disk read latency | Alert is triggered when all of the symptom sets are true.<br><br>Symptom set is met when any of the symptoms are true.<br>• Pod CPU swap wait is at Warning level<br>• Pod CPU swap wait is at Immediate level<br>• Pod CPU swap wait is at Critical level<br><br>Symptom set is met when all of the symptoms are true.<br>• Pod disk read latency at Warning level<br>• VMware Tools is running<br>• Pod does not have memory ballooning | |
| Pod has memory contention due to memory compression, ballooning, or swapping | Alert is triggered when all of the symptom sets are true:<br>• Pod memory limit is set<br><br>Symptom set is met when any of the symptoms are true.<br>• Pod memory contention is at Critical level<br>• Pod memory contention is at Immediate level<br>• Pod memory contention is at warning level<br>• Pod memory is compressed<br>• Pod memory ballooning is at Warning level<br>• Pod memory ballooning is at Immediate level<br>• Pod memory ballooning is at Critical level<br>• Pod is using swap | |
| Pod is demanding more CPU than the configured limit | Symptom set is met when all of the symptoms are true.<br>• Pod CPU limit is set<br>• CPU Demand is greater than configured limit | |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Pod is experiencing memory compression, ballooning, or swapping due to memory limit | Alert is triggered when all of the symptom sets are true.<br>• Pod memory limit is set<br>• Pod memory demand exceeds configured memory limit<br><br>Symptom set is met whenanyof the symptoms are true.<br>• Pod memory is compressed<br>• Pod memory ballooning is at Warning level<br>• Pod memory ballooning is at Immediate level<br>• Pod memory ballooning is at Critical level<br>• Pod is using swap | |
| Pod is in an invalid or orphaned state | Symptom set is met when any of the symptoms are true.<br>• Pod is in invalid state<br>• Pod is orphaned | |
| Pod on a host with BIOS power management not set to OS controlled is facing CPU contention | Alert is triggered when all of the symptom sets are true:<br>• Pod CPU contention at critical level<br><br>Symptom set is true when all of parent host system exhibit the following symptom.<br>• Host power management technology is not set to OS Controlled | |
| Pod on a host with BIOS power management not set to OS controlled is facing CPU contention | Alert is triggered when all of the symptom sets are true.<br><br>Symptom set is met when all of the symptoms are true.<br>• Pod CPU contention is elevated<br>• Pod CPU contention at critical level<br><br>Symptom set is true whenallof parent host system exhibit the following symptom.<br>• Host power management technology is not set to OS Controlled | |
| Pod on a host with BIOS power management set to OS controlled is facing CPU contention | Alert is triggered when all of the symptom sets are true. | |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | Symptom set is met when all of the symptoms are true.<br>• Pod CPU contention is elevated<br>• Pod CPU contention is elevated<br><br>Symptom set is true when all of parent host system exhibit the following symptom.<br>• Host power management technology is not set to OS Controlled | |
| Pod on a host with BIOS power management set to OS controlled is facing CPU contention | Alert is triggered when all of the symptom sets are true.<br><br>Symptom set is met when any of the symptoms are true.<br>• Pod CPU contention is elevated<br>• Pod CPU contention is elevated<br>• Pod CPU contention at critical level<br><br>Symptom set is true when all of parent host system exhibit the following symptom.<br>• Host power management technology is not set to OS Controlled | |
| vSphere HA failed to restart a network isolated Pod | vSphere HA failed to restart a network isolated Pod | |

### VMware Cloud on AWS Alert Definitions

Alert definitions are combinations of symptoms and recommendations that identify problem areas in your environment and generate alerts on which you can act. Symptom and alert definitions are defined for **VMware Cloud on AWS** objects.

#### Health/Symptom-Based

These alert definitions have the following impact and criticality information.

**Impact**

Risk

**Criticality**

Symptom-based

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| Number of SDDCs in this organization is exceeding the supported configuration maximum | VMC Configuration Maximum limits breached. The number of SDDCs in | • Please refer to VMC on AWS guide listed here. |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | this organization is over the supported limit. | • A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings. If the Soft Limit is already increased by VMware Support, and it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |
| Number of hosts per SDDC is at the supported configuration maximum | VMC Configuration Maximum limits are maxed out. The number of hosts in this SDDC is at the supported limit. | Please refer to VMware Cloud on AWS Configuration Maximum guide. |
| Number of clusters per SDDC soft limit is exceeding the supported configuration maximum | VMC Configuration Maximum limits are maxed out. Maximum number of clusters soft limit is over the supported limit. | • Please refer to VMware Cloud on AWS Configuration Maximum guide.<br>• A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |
| Number of virtual machines per SDDC is at the supported configuration maximum | VMC Configuration Maximum limits are maxed out. Number of virtual machines per SDDC is at the supported maximum | Please refer to VMware Cloud on AWS Configuration Maximum guide. |
| Number of linked VPCs in this SDDC is at the supported configuration maximum | VMC Configuration Maximum limits are maxed out. The number of linked VPCs in this SDDC is at the supported limit. | Please refer to VMC on AWS guide listed here. |
| Number of SDDCs in this organization is at the supported configuration maximum | VMC Configuration Maximum limits are maxed out. The number of SDDCs in this organization is at the supported limit. | • Please refer to VMC on AWS guide listed here.<br>• A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |
| Number of Public IP Addresses (Elastic IPs) per organization is exceeding the supported configuration maximum | VMC Configuration Maximum limits breached. The Number of Public IP Addresses (Elastic IPs) per organization is over the supported limit. | • Please refer to VMC on AWS guide listed here.<br>• A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |
| Number of clusters per SDDC hard limit is at supported configuration maximum | VMC Configuration Maximum limits are maxed out. Maximum number of clusters hard limit is at supported configuration maximum | Please refer to VMware Cloud on AWS Configuration Maximum guide. |
| Number of virtual machines per SDDC is exceeding the supported configuration maximum | VMC Configuration Maximum limits are breached. The number of virtual machines per SDDC is exceeding the supported maximum | Please refer to VMware Cloud on AWS Configuration Maximum guide. |
| Number of linked VPCs in this SDDC is exceeding the supported configuration maximum | VMC Configuration Maximum limits are maxed out. The number of linked VPCs in this SDDC is over the supported limit. | Please refer to VMC on AWS guide listed here. |
| Number of clusters per SDDC hard limit is exceeding the supported configuration maximum | VMC Configuration Maximum limits are maxed out. Maximum number of clusters hard limit is over the supported limit. | Please refer to VMware Cloud on AWS Configuration Maximum guide. |
| Number of clusters per SDDC soft limit is at supported configuration maximum | VMC Configuration Maximum limits are maxed out. Maximum number of clusters soft limit is at supported configuration maximum | • Please refer to VMware Cloud on AWS Configuration Maximum guide.<br>• A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |
| Number of hosts per organization is exceeding the supported configuration maximum | VMC Configuration Maximum limits are breached. The number of hosts in this organization is over the supported limit. | • Please refer to VMC on AWS guide listed here.<br>• A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |
| Number of hosts per organization is at the supported configuration maximum | VMC Configuration Maximum limits are maxed out. The number of hosts in this organization is at the supported limit. | • Please refer to VMC on AWS guide listed here.<br>• A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and |

*Table continued on next page*

*Continued from previous page*

| Alert Definition | Symptoms | Recommendations |
|---|---|---|
| | | it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |
| Number of hosts per SDDC is exceeding the supported configuration maximum | VMC Configuration Maximum limits are breached. The number of hosts in this SDDC is over the supported limit. | Please refer to VMware Cloud on AWS Configuration Maximum guide. |
| Number of Public IP Addresses (Elastic IPs) per organization is at the supported configuration maximum | VMC Configuration Maximum limits are maxed out. The Number of Public IP Addresses (Elastic IPs) per organization is at the supported limit. | • Please refer to VMC on AWS guide listed here.<br>• A Soft Limit can be increased in certain cases. To know more about this soft limit please contact Support Offerings If the Soft Limit is already increased by VMware Support, and it is not reflected in VMware Cloud Foundation Operations automatically, then refer to the KB article, KB 2059936. |

### Alerts in VMware Infrastructure Health

The following alerts are triggered when any of the monitoring resources in VMware Infrastructure Health display an unexpected behavior.

**Table 524: Alerts in VMware Infrastructure Health**

| Alert | Description |
|---|---|
| vCenter app health is affected | This Alert is triggered when the health of the vCenter application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down<br>• Certificate expired<br>• Password expired<br>• Any vCenter service is down<br>• NTP Server is down<br>• Back Job Service is failing |
| NSX app health is affected | This Alert is triggered when the health of the NSX application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down<br>• Certificate expired<br>• Password expired<br>• Any NSX-T service is down |

*Table continued on next page*

*Continued from previous page*

| Alert | Description |
|---|---|
| SDDC Manager app health is affected | This Alert is triggered when the health of the SDDC Manager application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down<br>• Certificate expired<br>• Password expired |
| VIDM app health is affected | This Alert is triggered when the health of the VMware Identity Manager application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down |
| Automation app health is affected | This Alert is triggered when the health of the VMware Aria Automation application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down<br>• Any Automation service is down<br>• Event occurred on Automation app |
| Logs app health is affected | This Alert is triggered when the health of the VMware Cloud Foundation Operations for logs application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down |
| Orchestrator app health is affected | This Alert is triggered when the health of the VMware Aria Operations Management Pack for VMware Aria Automation Orchestrator application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down |
| SRM app health is affected | This Alert is triggered when the health of the VMware Site Recovery Manager application is affected. Any of the following symptoms can cause the health degradation.<br>• Network connectivity issue<br>• Application is down |
| Lifecycle Manager App health is affected | This Alert is triggered when the health of the Fleet Management application is affected. Any of the following symptoms can cause the health degradation.<br>• Application is down |
| Networks app health is affected | This Alert is triggered when the health of the VMware Aria Operations for Networks application is affected. Any of the following symptoms can cause the health degradation.<br>• Networks node CPU usage is high<br>• Networks node Disk usage is high<br>• Networks node RAM usage is high |

*Table continued on next page*

*Continued from previous page*

| Alert | Description |
|---|---|
| | • Networks platform node is not healthy<br>• Networks collector node is not healthy<br>• Networks service is down<br>• Networks grid usage is high<br>• Networks application not reachable<br>• Networks processing lag is high<br>• Networks indexer lag is high |

### Alerts in NSX

Alert definitions are combinations of symptoms and recommendations that identify problem areas in your environment and generate alerts on which you can act. You can import all NSX alerts into VMware Cloud Foundation Operations, allowing the users to manage and monitor these alerts from a single location. This integration between NSX and VMware Cloud Foundation Operations improves overall visibility and simplifies operations, making it easier to identify and address issues.

You can directly import alerts from VMware NSX (formerly known as VMware NSX-T) into VMware Cloud Foundation Operations by setting the **Import Alerts from NSX** field to **True** while configuring the NSX adapter instance. For details, see the topic 'Configuring the NSX Adapter' in the *VMware Aria Operations Configuration Guide*.

Depending on the number of alerts and the system's performance, importing the alerts into VMware Cloud Foundation Operations may take some time.

**NOTE**
Importing alerts is not supported for NSX Integration through VMware Cloud on AWS.

#### Existing Out-of-the-Box NSX Alerts

| Alert | Description |
|---|---|
| BFD service is disabled | Triggered when the BFD service is not enabled on the logical router. |
| Controller Cluster Status is not stable | Triggered when all the controller nodes are down in NSX. |
| Controller Node Connectivity is broken | Triggered when the controller node connection status is down in NSX. |
| ECMP service is disabled for Logical Router | Triggered when the ECMP service is not enabled on the logical router. |
| File System usage is higher than 70 percent | Triggered when the guest file systems usage of the Controller Virtual Machine is more than 70 percent. |
| File System usage is more than 75 percent | Triggered when the guest file systems usage of the Controller Virtual Machine is more than 75 percent. |
| File System usage is more than 85 percent | Triggered when the guest file systems usage of the Controller Virtual Machine is more than 85 percent. |
| File System usage of manager node is more than 70 percent | Triggered when the guest file systems usage of the manager node is more than 70 percent. |
| File System usage of manager node is more than 75 percent | Triggered when the guest file systems usage of the manager node is more than 75 percent. |

*Table continued on next page*

*Continued from previous page*

| Alert | Description |
|---|---|
| File System usage of manager node is more than 85 percent | Triggered when the guest file systems usage of the manager node is more than 85 percent. |
| Less than 3 controller nodes are deployed | Triggered when the NSX server has less than three controller nodes. |
| Less than 3 manager nodes are deployed | Triggered when NSX server has less than 3 manager nodes deployed. |
| Load Balancer Service operational status down | Triggered when the operational status of the load balancer service is down. |
| Load balancer service operational status error | Triggered when the operational status of the load balancer service contains error. |
| Load Balancer virtual server operational state detached | Triggered when the operational state of the load balancer virtual server is detached |
| Load Balancer virtual server operational state down | Triggered when the operational state of the load balancer virtual server is down |
| Logical Switch State has failed | Triggered when the state of logical switch has failed. |
| Logical Switch's admin state is not UP | Triggered when the admin state is disabled on the logical switch. |
| Management cluster's management status is not stable | Triggered when the management status of a management cluster is not stable. |
| Management service monitor runtime state has failed | Triggered when the monitor runtime state of the management service stops running. |
| Management Status is not stable | Triggered when the management status of any node on the controller cluster is down. |
| Manager node connectivity is broken | Triggered when the manager connection status of manager node is down. |
| NAT rule not configured | Triggered when the NAT rule on the logical router is not configured. |
| NSX Management service has failed | Triggered when the runtime state of management service on the NSX host is not running. |
| NTP Service is violating VMware NSX Security Configuration Guide | NTP Service is violating one or more guidelines in VMware NSX Security Configuration Guide |
| Route Advertisement service is disabled | Triggered when the route advertisement service is not enabled on the logical router. |
| Route Redistribution service is disabled | Triggered when the route redistribution service is not enabled on the logical router. |
| SSH Service is violating VMware NSX Security Configuration Guide | SSH Service is violating one or more guidelines in VMware NSX Security Configuration Guide |
| Static Route not configured | Triggered when the static route on the logical router is not configured. |
| Transport Host node is in Failed/Error state | Triggered when the host node in NSX is in error or failed state due to one of the following reasons:<br>• Edge configuration error<br>• Installation failure<br>• Uninstallation failure<br>• Upgrade failure<br>• Virtual Machine deployment failure |

*Table continued on next page*

*Continued from previous page*

| Alert | Description |
|---|---|
| | • Virtual Machine power off failure<br>• Virtual Machine power on failure<br>• Virtual Machine undeployment failure |
| Transport node configuration state has failed | Triggered when the configuration state of transport node has failed. |
| Transport Node Controller/Manager Connectivity is not UP | Triggered when the transport node connectivity status is down in NSX. |

For details on VMware NSX alerts that can be imported into VMware Cloud Foundation Operations, refer to the VMware NSX Event Catalog.

## Property Definitions in VMware Aria OperationsVMware Cloud Foundation Operations

Properties are attributes of objects in the VMware Aria OperationsVMware Cloud Foundation Operations environment. You use properties in symptom definitions. You can also use properties in dashboards, views, and reports.

VMware Aria OperationsVMware Cloud Foundation Operations uses adapters to collect properties for target objects in your environment. Property definitions for all objects connected through the vCenter adapter are provided. The properties collected depend on the objects in your environment.

You can add symptoms based on properties to an alert definition so that you are notified if a change occurs to properties on your monitored objects. For example, disk space is a hardware property of a virtual machine. You can use disk space to define a symptom that warns you when the value falls below a certain numeric value.

You can add symptoms based on properties to an alert definition so that you are notified if a change occurs to properties on your monitored objects. For example, disk space is a hardware property of a virtual machine. You can use disk space to define a symptom that warns you when the value falls below a certain numeric value. See the *VMware Aria OperationsVMware Cloud Foundation Operations User Guide*.

VMware Aria OperationsVMware Cloud Foundation Operations generates Object Type Classification and Subclassification properties for every object. You can use object type classification properties to identify whether an object is an adapter instance, custom group, application, tier, or a general object with property values *ADAPTER_INSTANCE*, *GROUP*, *BUSINESS_SERVICE*, *TIER*, or *GENERAL*, respectively.

### Properties for vCenter Server Components

The VMware vSphere solution is installed with VMware Aria OperationsVMware Cloud Foundation Operationsand includes the vCenter adapter. VMware Aria OperationsVMware Cloud Foundation Operations uses the vCenter adapter to collect properties for objects in the vCenter system.

vCenter components are listed in the `describe.xml` file for the vCenter adapter. The following example shows the runtime property `memoryCap` or Memory Capacity for the virtual machine in the `describe.xml`.

```
<ResourceGroup instanced="false" key="runtime" nameKey="5300" validation="">
   <ResourceAttribute key="memoryCap" nameKey="1780" dashboardOrder="200" dataType="float"

                       defaultMonitored="true" isDiscrete="false" isRate="false" maxVal=""

                       minVal="" isProperty="true" unit="kb"/>
</ResourceGroup>
```

The `ResourceAttribute` element includes the name of the property that appears in the UI and is documented as a Property Key. `isProperty = "true"` indicates that `ResourceAttribute` is a property.

### vCenter Server Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects summary and event properties for system objects.

**Table 525: Summary Properties Collected for vCenterSystem Objects**

| Property Key | Property Name | Description |
|---|---|---|
| summary\|version | Version | Version |
| summary\|vcuuid | VirtualCenter ID | Virtual Center ID |
| summary\|vcfullname | Product Name | Product Name |

**Table 526: Event Properties Collected for vCenterSystem Objects**

| Property Key | Property Name | Description |
|---|---|---|
| event\|time | Last VC Event Time | Last Virtual Center Event Time |
| event\|key | Last VC Event ID | Last Virtual Center Event ID |

**Table 527: Custom Field Manager Property Collected for vCenterSystem Objects**

| Property Key | Property Name | Description |
|---|---|---|
| CustomFieldManager\|CustomFieldDef | Custom Field Def | Custom Field Def for vCenter Tagging information at the Adapter level. |

**Table 528: Compliance Configuration Related Properties for vCenterSystem Objects**

| Property Key | Property Name | Description |
|---|---|---|
| vc_appliance\|hasAccessSSH | Appliance\|Access SSH | Access SSH |
| vc_appliance\|networkNICs | Appliance\|Number of NICs | NICs |

### Virtual Machine Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, runtime, CPU, memory, network I/O, and properties about summary use for virtual machine objects. Properties are collected with the first cycle of data collection. Once collected, the next property collection occurs only when there is data change. In case of no data change, no property is collected.

**Table 529: Properties Collected for Virtual Machines with vSAN Adapter**

| Property Key | Property Name | Description |
|---|---|---|
| **NOTE**<br>The following Disk Space properties are displayed by the virtual machine object only when the vSAN adapter is configured with vCenter. | | |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| Virtual Disk:scsi0:0\|Storage Policy\| Cache Reservation | Storage Policy\| Cache Reservation | This property helps you with flash read cache reservation. |
| Virtual Disk:scsi0:0\|Storage Policy\| Checksum Disabled | Storage Policy\| Checksum Disabled | This property checks whether checksum is disabled for the disk object. |
| Virtual Disk:scsi0:0\|Storage Policy\| Encryption Service | Storage Policy\| Encryption Service | This property checks whether the encryption service is used or not. |
| Virtual Disk:scsi0:0\|Storage Policy\| Failures to Tolerate | Storage Policy\| Failures to Tolerate | This property gives the number of host or disk failures the storage object can tolerate. |
| Virtual Disk:scsi0:0\|Storage Policy\| Force Provisioning | Storage Policy\| Force Provisioning | This property checks whether force provisioning is disabled for the disk object. |
| Virtual Disk:scsi0:0\|Storage Policy\| Object Space Reservation | Storage Policy\| Object Space Reservation | This property helps you to check the object space reservation. |
| Virtual Disk:scsi0:0\|Storage Policy\| Replica Preference | Storage Policy\| Replica Preference | This property helps you to identify the fault tolerance method. |
| Virtual Disk:scsi0:0\|Storage Policy\|Site Disaster Tolerance | Storage Policy\|Site Disaster Tolerance | This property gives the number of fault domain failures, the storage object can tolerate. |
| Virtual Disk:scsi0:0\|Storage Policy\| Space Efficiency | Storage Policy\| Space Efficiency | This property helps you with the space efficiency method. |
| Virtual Disk:scsi0:0\|Storage Policy\| Storage Policy in Use | Storage Policy in Use | This property gives the storage policy associated with the virtual disk. |
| Virtual Disk:scsi0:0\|Storage Policy\| Stripe Width | Storage Policy\| Stripe Width | This property gives the number of disk stripes per object. |

**Table 530: Properties Collected for System**

| Property Key | Property Name | Description |
|---|---|---|
| system\|notes | System\|Notes | This property helps you to track the System notes defined in vCenter. |

**Table 531: VMware Aria Automation Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| vRealize Automation\|Blueprint Name | Blueprint Name | Virtual machines deployed by VMware Aria Automation to be excluded from workload placements. |

**Table 532: Properties Collected for Virtual Machine Objects to Support VIN Adapter Localization**

| Property Key | Property Name | Description |
|---|---|---|
| RunsOnApplicationComponents | Application components running on the Virtual Machine | Application components running on the Virtual Machine |
| DependsOnApplicationComponents | Application components the Virtual Machine depends on | Application components running on other machines that this Virtual Machine depends on. |

**Table 533: Properties Collected for Guest File Systems**

| Property Key | Property Name | Description |
|---|---|---|
| guestfilesystem\|capacity_property | Guest File System stats\| Guest File System Capacity Property | This property is disabled by default. |
| guestfilesystem\|capacity_property_total | Guest File System stats\|Total Guest File System Capacity Property(gb) | This property is disabled by default. |

**Table 534: Properties Collected for Disk Space Objects**

| Property Key | Property Name | Description |
|---|---|---|
| diskspace\|snapshot\|creator | Disk Space\|Snapshot\| Creator | This property is disabled by default. |
| diskspace\|snapshot\|description | Disk Space\|Snapshot\| Description | This property is disabled by default. |

**Table 535: Configuration Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|ctkEnabled | Configuration\|Changed Block Tracking | This property displays if Change Block Tracking is enabled, if enabled, tracks changes on disk sectors. This helps in performing incremental backups on the VM. |
| Configuration\|Hardware\|Number of CPUs (vCPUs) | Number of CPUs (vCPUs) | This property displays the number of CPUs configured in the VM, the count includes both in the vSocket and the vCore. |
| Configuration\|Number of RDMs | Number of RDMs | This property displays the number of RDMs configured in the VM. This property is enabled by default. |
| Configuration\|Number of Virtual Disks | Number of Virtual Disks | This property displays the number of virtual disks configured in the VM, the count includes the RDMs. |
| Configuration\|Hardware\|Number of VMDKs | Number of VMDKs | This property displays the number of VMDKs configured in the VM. This property is enabled by default. |
| Summary\|Is Horizon Managed | Is Horizon Managed | This property displays whether the selected object is managed by Horizon or not. |
| summary\|datastoreClusters | Summary\|Datastore Cluster(s) | This property is applicable only if the VM belongs to a datastore cluster.<br>**NOTE**<br>A VM with multiple virtual disks can belong to multiple datastore clusters. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| Summary\|Configuration\|Number of NICs | Number of NICs | his property displays the number of NICs configured in the VM. |
| config\|name | Name | This property displays the virtual object name |
| config\|guestFullName | Guest OS from vCenter | This property is set by the vCenter during the VM creation. It may differ from the value of the Guest/ |
| config\|hardware\|numCpu | Number of virtual CPUs | Number of virtual CPUs |
| config\|hardware\|memoryKB | Memory | Memory |
| config\|hardware\|thinEnabled | Thin Provisioned Disk | Indicates whether thin provisioning is enabled |
| config\|hardware\|diskSpace | Disk Space | Disk Space |
| config\|cpuAllocation\|reservation | Reservation | CPU reservation |
| config\|cpuAllocation\|limit | Limit | CPU limit |
| config\|cpuAllocation\|shares\|shares | Shares | CPU shares |
| config\|memoryAllocation\|reservation | Reservation | CPU reservation |
| config\|memoryAllocation\|limit | Limit | Limit |
| config\|memoryAllocation\|shares\|shares | Shares | Memory shares |
| config\|extraConfig\|mem_hotadd | Memory Hot Add | Memory Hot Add Configuration |
| config\|extraConfig\|vcpu_hotadd | VCPU Hot Add | VCPU Hot Add Configuration |
| config\|extraConfig\|vcpu_hotremove | VCPU Hot Remove | VCPU Hot Remove Configuration |
| config\|security\|disable_autoinstall | Disable tools auto install (isolation.tools.autoInstall.disable) | Disable tools auto install (isolation.tools.autoInstall.disable) |
| config\|security\|disable_console_copy | Disable console copy operations (isolation.tools.copy.disable) | Disable console copy operations (isolation.tools.copy.disable) |
| config\|security\|disable_console_dnd | Disable console drag and drop operations (isolation.tools.dnd.disable) | Disable console drag and drop operations (isolation.tools.dnd.disable) |
| config\|security\| enable_console_gui_options | Enable console GUI operations (isolation.tools.setGUIOptions.enable) | Enable console GUI operations (isolation.tools.setGUIOptions.enable) |
| config\|security\|disable_console_paste | Disable console paste operations (isolation.tools.paste.disable) | Disable console paste operations (isolation.tools.paste.disable) |
| config\|security\| disable_disk_shrinking_shrink | Disable virtual disk shrink (isolation.tools.diskShrink.disable) | Disable virtual disk shrink (isolation.tools.diskShrink.disable) |
| config\|security\| disable_disk_shrinking_wiper | Disable virtual disk wiper (isolation.tools.diskWiper.disable) | Disable virtual disk wiper (isolation.tools.diskWiper.disable) |
| config\|security\|disable_hgfs | Disable HGFS file transfers (isolation.tools.hgfsServerSet.disable) | Disable HGFS file transfers (isolation.tools.hgfsServerSet.disable) |
| config\|security\| disable_independent_nonpersistent | Avoid using independent nonpersistent disks (scsiX:Y.mode) | Avoid using independent nonpersistent disks (scsiX:Y.mode) |
| config\|security\|enable_intervm_vmci | Enable VM-to-VM communication through VMCI (vmci0.unrestricted) | Enable VM-to-VM communication through VMCI (vmci0.unrestricted) |
| config\|security\|enable_logging | Enable VM logging (logging) | Enable VM logging (logging) |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| config\|security\|disable_monitor_control | Disable VM Monitor Control (isolation.monitor.control.disable) | Disable VM Monitor Control (isolation.monitor.control.disable) |
| config\|security\| enable_non_essential_3D_features | Enable 3D features on Server and desktop virtual machines (mks.enable3d) | Enable 3D features on Server and desktop virtual machines (mks.enable3d) |
| config\|security\| disable_unexposed_features_autologon | Disable unexposed features - autologon (isolation.tools.ghi.autologon.disable) | Disable unexposed features - autologon (isolation.tools.ghi.autologon.disable) |
| config\|security\| disable_unexposed_features_biosbbs | Disable unexposed features - biosbbs (isolation.bios.bbs.disable) | Disable unexposed features - biosbbs (isolation.bios.bbs.disable) |
| config\|security\| disable_unexposed_features_getcreds | Disable unexposed features - getcreds (isolation.tools.getCreds.disable) | Disable unexposed features - getcreds (isolation.tools.getCreds.disable) |
| config\|security\| disable_unexposed_features_launchmenu | Disable unexposed features - launchmenu (isolation.tools.ghi.launchmenu.change) | Disable unexposed features - launchmenu (isolation.tools.ghi.launchmenu.change) |
| config\|security\| disable_unexposed_features_memsfss | Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable) | Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable) |
| config\|security\| disable_unexposed_features_protocolhandler | Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable) | Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable) |
| config\|security\| disable_unexposed_features_shellaction | Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable) | Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable) |
| config\|security\| disable_unexposed_features_toporequest | Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable) | Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable) |
| config\|security\| disable_unexposed_features_trashfolderstate | Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable) | Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable) |
| config\|security\| disable_unexposed_features_trayicon | Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable) | Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable) |
| config\|security\| disable_unexposed_features_unity | Disable unexposed features - unity (isolation.tools.unity.disable) | Disable unexposed features - unity (isolation.tools.unity.disable) |
| config\|security\| disable_unexposed_features_unity_interlock | Disable unexposed features - unity-interlock (isolation.tools.unityInterlockOperation.disable) | Disable unexposed features - unity-interlock (isolation.tools.unityInterlockOperation.disable) |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| config\|security\|disable_unexposed_features_unity_taskbar | Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar.disable) | Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar.disable) |
| config\|security\|disable_unexposed_features_unity_unityactive | Disable unexposed features - unity-unityactive (isolation.tools.unityActive.disable) | Disable unexposed features - unity-unityactive (isolation.tools.unityActive.disable) |
| config\|security\|disable_unexposed_features_unity_windowcontents | Disable unexposed features - unity-windowcontents (isolation.tools.unity.windowContents.disable) | Disable unexposed features - unity-windowcontents (isolation.tools.unity.windowContents.disable) |
| config\|security\|disable_unexposed_features_unitypush | Disable unexposed features - unitypush (isolation.tools.unity.push.update.disable) | Disable unexposed features - unitypush (isolation.tools.unity.push.update.disable) |
| config\|security\|disable_unexposed_features_versionget | Disable unexposed features - versionget (isolation.tools.vmxDnDVersionGet.disable) | Disable unexposed features - versionget (isolation.tools.vmxDnDVersionGet.disable) |
| config\|security\|disable_unexposed_features_versionset | Disable unexposed features - versionset (solation.tools.guestDnDVersionSet.disable) | Disable unexposed features - versionset (solation.tools.guestDnDVersionSet.disable) |
| config\|security\|disable_vix_messages | Disable VIX messages from the VM (isolation.tools.vixMessage.disable) | Disable VIX messages from the VM (isolation.tools.vixMessage.disable) |
| config\|security\|enable_vga_only_mode | Disable all but VGA mode on virtual machines (svga.vgaOnly) | Disable all but VGA mode on virtual machines (svga.vgaOnly) |
| config\|security\|limit_console_connection | Limit number of console connections (RemoteDisplay.maxConnection) | Limit number of console connections (RemoteDisplay.maxConnection) |
| config\|security\|limit_log_number | Limit number of log files (log.keepOld) | Limit number of log files (log.keepOld) |
| config\|security\|limit_log_size | Limit log file size (log.rotateSize) | Limit log file size (log.rotateSize) |
| config\|security\|limit_setinfo_size | Limit VMX file size (tools.setInfo.sizeLimit) | Limit VMX file size (tools.setInfo.sizeLimit) |
| config\|security\|enable_console_VNC | Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled) | Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled) |
| config\|security\|disable_device_interaction_connect | Disable unauthorized removal, connection of devices (isolation.device.connectable.disable) | Disable unauthorized removal, connection of devices (isolation.device.connectable.disable) |
| config\|security\|disable_device_interaction_edit | Disable unauthorized modification of devices (isolation.device.edit.disable) | Disable unauthorized modification of devices (isolation.device.edit.disable) |
| config\|security\|enable_host_info | Enable send host information to guests (tools.guestlib.enableHostInfo) | Enable send host information to guests (tools.guestlib.enableHostInfo) |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| config\|security\|network_filter_enable | Enable dvfilter network APIs (ethernetX.filterY.name) | Enable dvfilter network APIs (ethernetX.filterY.name) |
| config\|security\| vmsafe_cpumem_agentaddress | VMsafe CPU/memory APIs - IP address (vmsafe.agentAddress) | VMsafe CPU/memory APIs - IP address (vmsafe.agentAddress) |
| config\|security\|vmsafe_cpumem_agentport | VMsafe CPU/memory APIs - port number (vmsafe.agentPort) | VMsafe CPU/memory APIs - port number (vmsafe.agentPort) |
| config\|security\|vmsafe_cpumem_enable | Enable VMsafe CPU/memory APIs (vmsafe.enable) | Enable VMsafe CPU/memory APIs (vmsafe.enable) |
| config\|security\|disconnect_devices_floppy | Disconnect floppy drive | Disconnect floppy drive |
| config\|security\|disconnect_devices_cd | Disconnect CD-ROM | Disconnect CD-ROM |
| config\|security\|disconnect_devices_usb | Disconnect USB controller | Disconnect USB controller |
| config\|security\| disconnect_devices_parallel | Disconnect parallel port | Disconnect parallel port |
| config\|security\|disconnect_devices_serial | Disconnect serial port | Disconnect serial port |
| config\|faultTolerant | config\|faultTolerant | |

**NOTE**
Security properties not collected by default. They are collected only if the *vSphere Hardening Guide* policy is applied to the objects, or if the *vSphere Hardening Guide* alerts are manually enabled in the currently applied policy.

**Table 536: Runtime Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| runtime\|memoryCap | Memory Capacity | Memory Capacity |

**Table 537: CPU Usage Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| cpu\|limit | CPU limit | CPU limit |
| cpu\|reservation | CPU reservation | CPU reservation |
| cpu\|speed | CPU | CPU Speed |

**Table 538: Memory Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| mem\|host_limit | VM Limit | Mem Machine Limit |
| mem\|host_reservation | Memory\|VM Reservation(kb) | This property is disabled by default. |

**Table 539: Network Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| net:<vnic_id>\|portGroup | Network:<vnic_id>\|Port Group | This is a property at each virtual NIC level, not at the VM level. A VM with multiple NIC connects to different port groups, with this property you can identify to which port group the vnic blongs and has mapping with. |
| net:<nic_key>\| uptCompatibilityEnabled | Network:<nic_key>\|Direct Path I/O Status | If there is the VM, with cofigured direct path IO - IO directly comes to VM, bypassing the hypervisor. The property shows the status. |
| net\|mac_address | Mac Address | Mac Address |
| net\|ip_address | IP Address | IP Address |
| net\|vnic_label | Network:<ID>\|Label | This property is disabled by default. |
| net\|nvp_vm_uuid | Network I/O\|NVP VM UUID | This property is disabled by default. |
| net\|vnic_type | Network I/O\|Virtual NIC Type | This property is disabled by default. |
| net\|ipv6_address | Network\|IPv6 Address | This property is disabled by default. |
| net\|ipv6_prefix_length | Network\|IPv6 Prefix Length | This property is disabled by default. |
| net\|default_gateway | Network\|Network I/O\|Default Gateway | This property is disabled by default. |
| net\|subnet_mask | Network\|Subnet Mask | This property is disabled by default. |

**Table 540: Summary Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| summary\|customTag\|customTagValue | Value | Custom Tag Value |
| summary\|tag | vSphere Tag | vSphere Tag Name |
| summary\|parentCluster | Parent Cluster | Parent Cluster |
| summary\|parentHost | Parent Host | Parent Host |
| summary\|parentDatacenter | Parent data center | Parent data center |
| summary\|parentVcenter | Parent vCenter | Parent vCenter |
| summary\|guest\|fullName | Guest OS Full Name | This property is provided by the VMware Tools. It will differ to the value set in vCenter if the Guest OS was upgraded, or if a different Guest OS was installed. |
| summary\|guest\|ipAddress | Guest OS IP Address | Guest OS IP Address |
| summary\|guest\|toolsRunningStatus | Tools Running Status | Guest Tools Running Status |
| summary\|guest\|toolsVersionStatus2 | Tools Version Status | Guest Tools Version Status 2 |
| summary\|guest\| vrealize_operations_agent_id | vRealize Operations Agent ID | An ID to identify a VM in Agent Adapter's world. |
| summary\|guest\| vrealize_operations_euc_agent_id | vRealize Operations Euc Agent ID | An ID to identify a VM in Agent Adapter's world. |
| summary\|config\|numEthernetCards | Number of NICs | Number of NICs |
| summary\|config\|isTemplate | VM Template | Indicates whether it is a VM Template. |
| summary\|runtime\|powerState | Power State | Power State |
| summary\|runtime\|connectionState | Connection State | Connection State |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| summary\|config\|appliance | Appliance | Appliance |
| summary\|config\|productName | Product Name | Product Name |
| summary\|UUID | UUID | Unique UUID instance in vCenter that identifies all the virtual machine instances. |
| summary\|smbiosUUID | SMBIOS UUID | System Management BIOS UUID of a virtual machine. |

**Table 541: Virtual Disk Properties Collected for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| virtualDisk:<scsi_id>\|encryptionStatus | Virtual Disk:<scsi_id>\|Encryption Status | This property provides encryption status on SCSI on VDMK layer. The property is per virtual disk and does not apply to RDM disk. |
| virtualDisk:<scsi_controller>\|iopsLimit | Virtual Disk:<scsi_controller>\|IOPS Limit | This property shows the set IOPS limit on virtual disk on SCSI level, when IOPS limit is set beyond the disk. This helps you to understand the current performance state, why the IOPS is not increasing and limited. |
| Virtual Disk:scsi0:0\|Configuired Size(GB) | Configuired Size(GB) | This property displays the disk space configured for the virtual disk. |
| Virtual Disk:scsi0:0\|Datastore | Datastore | This property displays the name of the datastore where the scsi disk is present. If an RDM is present, only a pointer to the RDM is present in the datastore. |
| Virtual Disk:scsi0:0\|Disk Mode | Disk Mode | This property determines how a virtual disk is affected by snapshots. The disk mode acts on each individual VMDK, not on the whole VM. The available options are Independent persistent, Persistent, Non-persistent, this property is disabled by default. |
| Virtual Disk:scsi0:0\|SCSI Bus Sharing | SCSI Bus Sharing | This property sets the type of bus sharing for the VM and determines whether to share the bus or not. Depending on the type of bus sharing, the VM can access the same virtual disk simultaneously on the same server or any other server. The available options are None, Physical, Virtual, this property is disabled by default. |
| Virtual Disk:scsi0:0\|SCSI Controller Type | SCSI Controller Type | This virtual storage controller property connects the virtual and physical disks to the VM. The available options are LSI SAS/PVSCSI, this property is disabled by default. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| Virtual Disk:scsi0:0\|Virtual Disk Sharing | Virtual Disk Sharing | This property allows VMFS-backed disks to be shared by multiple VMs. The available options are Unspecified, No sharing, Multi-Writer. You can use this option to disable protection for certain cluster aware applications, where the applications ensures that simultaneous write operation from two VMs does not induce data loss. This property is disabled by default. |
| Virtual Disk:scsi0:0\|Virtual Device Node | Virtual Device Node | This property determines the virtual device bus location. The virtual disks are enumerated starting with the first controller. This property is disabled by default. |
| Virtual Disk:scsi0:0\|Is RDM | Is RDM | This property indicates whether the virtual disk is an RDM or not. This property is enabled by default. |
| Virtual Disk\|File Name | File Name | This property is disabled by default. |
| Virtual Disk\|Label | Label | This property displays the device label. |
| Virtual Disk:scsi1:1\|Compatibility Mode | Compatibility Mode | This property displays the compatibility mode for the RDMs. The options are physical and virtual. Virtual mode specifies the full virtualization of the mapped device whereas physical mode specifies minimal SCSI virtualization of the mapped device. This property is disabled by default in the base policy. |

**Table 542: Virtual Disk Properties Collected for POD Objects**

| Property Key | Property Name | Description |
|---|---|---|
| Virtual Disk:scsi0:0\|Virtual Device Node | Virtual Device Node | This property determines the virtual device bus location. The virtual disks are enumerated starting with the first controller. This property is disabled by default. |
| Virtual Disk:scsi0:0\|Virtual Disk Sharing | Virtual Disk Sharing | This property allows VMFS-backed disks to be shared by multiple VMs. The available options are Unspecified, No sharing, Multi-Writer. You can use this option to disable protection for certain cluster aware applications, where the applications ensures that simultaneous write operation from two PODs does not induce data loss. This property is disabled by default. |
| Virtual Disk:scsi0:0\|Disk Mode | Disk Mode | This property determines how a virtual disk is affected by snapshots. The disk mode acts on each individual VMDK, not on the whole POD. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| | | The available options are Independent persistent, Persistent, Non-persistent, this property is disabled by default. |
| Virtual Disk:scsi0:0\|SCSI Controller Type | SCSI Controller Type | This virtual storage controller property connects the virtual and physical disks to the POD. The available options are LSI SAS/ PVSCSI, this property is disabled by default. |

**Table 543: Datastore Properties Collected for Virtual Machine Properties**

| Property Key | Property Name | Description |
|---|---|---|
| datastore\|maxObservedNumberRead | Datastore I/O\|Highest Observed Number of Read Requests | |
| datastore\|maxObservedNumberWrite | Datastore I/O\|Highest Observed Number of Write Requests | |
| datastore\|maxObservedOIO | Datastore I/O\|Highest Observed Outstanding Requests | |
| datastore\|maxObservedRead | Datastore I/O\|Highest Observed Read Rate(kbps) | |
| datastore\|maxObservedWrite | Datastore I/O\|Highest Observed Write Rate(kbps) | |

**Table 544: Compliance Configuration Related Properties for Virtual Machine Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|security\| disconnect_devices_virtualhdaudiocard | Configuration\|Security\|Virtual HD Audio Card Disconnected | NIL |
| config\|security\| disconnect_devices_virtualahcicontroller | Configuration\|Security\|Virtual AHCI Controller Disconnected | NIL |
| config\|security\| disconnect_devices_virtualensoniq1371 | Configuration\|Security\|Virtual Ensoniq 1371 Disconnected | NIL |

Datastore properties collected for virtual machine objects have been disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

### Host System Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, hardware, runtime, CPU, network I/O, and properties about summary use for host system objects.

**Table 545: GPU properties collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| gpu\|<GPU-id>\|active_type | GPU\|<GPU-id>\|Active Type | Active Type |
| gpu\|<GPU-id>\|configured_type | GPU\|<GPU-id>\|Configured Type | Configured Type |
| gpu\|<GPU-id>\|device_name | GPU\|<GPU-id>\|Device Name | Device name |
| gpu\|<GPU-id>\|vendor_name | GPU\|<GPU-id>\|Vendor Name | Vendor name |
| gpu\|assignmentPolicy | GPU\|Assignment Policy | Assignment Policy |

**Table 546: Configuration Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|name | Name | Name |
| config\|diskSpace | Disk Space | Disk Space |
| config\|network\|nnic | Number of NICs | Number of NICs |
| config\|network\|linkspeed | Average Physical NIC Speed | Average Physical NIC Speed |
| config\|network\|dnsserver | DNS Server | List of DNS Servers |
| config\|product\|productLineId | Product Line ID | Product Line ID |
| config\|product\|apiVersion | API Version | API Version |
| config\|storageDevice\|plugStoreTopology\|numberofPath | Total number of Path | Total number of storage paths |
| config\|storageDevice\|multipathInfo\|numberofActivePath | Total number of Active Path | Total number of active storage paths |
| config\|storageDevice\|multipathInfo\|multipathPolicy | Multipath Policy | Multipath Policy |
| config\|hyperThread\|available | Available | Indicates whether hyperthreading is supported by the server |
| config\|hyperThread\|active | Active | Indicates whether hyperthreading is active |
| config\|ntp\|server | NTP Servers | NTP Servers |
| config\|security\|ntpServer | NTP server | NTP server |
| config\|security\|enable_ad_auth | Enable active directory authentication | Enable active directory authentication |
| config\|security\|enable_chap_auth | Enable mutual chap authentication | Enable mutual chap authentication |
| config\|security\|enable_auth_proxy | Enable authentication proxy (UserVars.ActiveDirectoryVerifyCAMCertificate) | Enable authentication proxy (UserVars.ActiveDirectoryVerifyCAMCertificate) |
| config\|security\|syslog_host | Remote log host (Syslog.global.logHost) | Remote log host (Syslog.global.logHost) |
| config\|security\|dcui_access | Users who can override lock down mode and access the DCUI (DCUI.Access) | Users who can override lock down mode and access the DCUI (DCUI.Access) |
| config\|security\|shell_interactive_timeout | Shell interactive timeout (UserVars.ESXiShellInteractiveTimeOut) | Shell interactive timeout (UserVars.ESXiShellInteractiveTimeOut) |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| config\|security\|shell_timeout | Shell timeout (UserVars.ESXiShellTimeOut) | Shell timeout (UserVars.ESXiShellTimeOut) |
| config\|security\|dvfilter_bind_address | Dvfilter bind ip address (Net.DVFilterBindIpAddress) | Dvfilter bind ip address (Net.DVFilterBindIpAddress) |
| config\|security\|syslog_dir | Log directory (Syslog.global.logDir) | Log directory (Syslog.global.logDir) |
| config\|security\|firewallRule\|allowedHosts | Allowed hosts | Allowed hosts in the firewall configuration |
| config\|security\|service\|isRunning | Running | Indicates whether a service is running or not. Services are: Direct Console UI, ESXi shell, SSH, or NTP Daemon. |
| config\|security\|service\|ruleSet | Ruleset | Ruleset for each service. |
| config\|security\|service\|policy | Policy | Policy for each service. |
| config\|security\|tlsdisabledprotocols | TLS Disabled Protocols | TLS Disabled Protocols |

**NOTE**
Security properties not collected by default. They are collected only if the *vSphere Hardening Guide* policy is applied to the objects, or if the *vSphere Hardening Guide* alerts are manually enabled in the currently applied policy.

**Table 547: Cost Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| Cost\|Energy Consumed (Joule) | Energy Consumed (Joule) | Displays the energy consumed in Joules. |
| Cost\|Number of Rack Units | Number of Rack Units | Displays the number of rack units in the host. |
| Cost\|OS Categories | OS Categories | Displays the operating system categories in the host. |
| Cost\|IsServerLeased | Is Server Leased | Displays whether the server is leased or not. |
| Cost\|RemainingDepreciationMonths | Remaining Depreciation Months | Displays the remaining number of depreciation months. |
| Cost\|ServerPurchaseCost | Server Purchase Cost | Server Purchase Cost is displayed in the currency format chosen. |
| Cost\|ServerPurchaseDate | Server Purchase Date | Server Purchase Date is displayed |

**Table 548: Hardware Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| hardware\|bioisReleaseDate | Hardware\|BIOS Release Date | This property displays the release date corresponding to the version of the installed BIOS. |
| hardware\|memorySize | Memory Size | Memory Size |
| hardware\|cpuInfo\|numCpuCores | Number of CPU Cores | Number of CPU Cores |
| hardware\|cpuInfo\|hz | CPU Speed per Core | CPU Speed per Core |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| hardware\|cpuInfo\|numCpuPackages | Number of CPU Packages | Number of CPU Packages |
| hardware\|cpuInfo\|<br>powerManagementPolicy | Active CPU Power Management Policy | Active CPU Power Management Policy |
| hardware\|cpuInfo\|<br>powerManagementTechnology | Power Management Technology | Power Management Technology |
| hardware\|cpuInfo\|biosVersion | BIOS Version | BIOS Version |
| hardware\|vendor | Hardware\|Vendor | Indicates the hardware manufacturer |

**Table 549: Runtime Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| runtime\|connectionState | Connection State | Connection State |
| runtime\|powerState | Power State | Power State |
| runtime\|maintenanceState | Maintenance State | Maintenance State |
| runtime\|memoryCap | Memory Capacity | Memory Capacity |

**Table 550: Configuration Manager Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| configManager\|memoryManager\|<br>consoleReservationInfo\|<br>serviceConsoleReserved | Service Console Reserved | Service console reserved memory |

**Table 551: CPU Usage Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| cpu\|speed | CPU | CPU Speed |
| cpu\|cpuModel | CPU Model | CPU Model |

**Table 552: Network Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| net:<pnic>\|configuredSpeed | Network:<pnic>\|Configured Speed | This property displays the configured network speed of the network card. If this is higher than actual, the card is not operating at full capacity. |
| net:<pnic>\|speed | Network:<pnic>\|Actual Speed | This property displays the actual operating speed of the network card, which can be lower than its configured capacity due to auto-negotiation. The options are Enabled or Disabled. |
| net\|maxObservedKBps | Highest Observed Throughput | Highest Observed Throughput (KBps) |
| net\|mgmt_address | Management Address | Management Address |
| net\|ip_address | IP Address | IP Address |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| net\|discoveryProtocol\|cdp\|managementIpAddress | Management IP Address | Management IP Address |
| net\|discoveryProtocol\|cdp\|systemName | System Name | System Name |
| net\|discoveryProtocol\|cdp\|portName | Port Name | Port Name |
| net\|discoveryProtocol\|cdp\|vlan | VLAN | VLAN |
| net\|discoveryProtocol\|cdp\|mtu | MTU | Maximum Transmission Unit |
| net\|discoveryProtocol\|cdp\|hardwarePlatform | Hardware Platform | Hardware Platform |
| net\|discoveryProtocol\|cdp\|softwareVersion | Software Version | Software Version |
| net\|discoveryProtocol\|lldp\|systemDescription | System Description | System Description |
| net\|discoveryProtocol\|lldp\|mtu | MTU | Maximum Transmission Unit |
| net\|discoveryProtocol\|lldp\|portDescription | Port Description | Port Description |
| net\|discoveryProtocol\|lldp\|aggregationStatus | Aggregation Status | Aggregation Status |
| net\|discoveryProtocol\|lldp\|managementIpAddress | Management IP Address | Management IP Address |
| net\|discoveryProtocol\|lldp\|systemName | System Name | System Name |
| net\|discoveryProtocol\|lldp\|portName | Port Name | Port Name |
| net\|discoveryProtocol\|lldp\|vlan | VLAN | VLAN |

**Table 553: System Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| sys\|build | Build number | VMWare build number |
| sys\|productString | Product String | VMWare product string |

**Table 554: Summary Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| Summary\|Is Horizon Managed | Is Horizon Managed | This property displays whether the selected object is managed by Horizon or not. |
| summary\|version | Version | This property displays theVersion. |
| summary\|hostuuid | Host UUID | This property displays the Host UUID. |
| summary\|evcMode | Current EVC Mode | This property displays the Current EVC Mode. |
| summary\|customTag\|customTagValue | Value | This property displays the Custom Tag Value. |
| summary\|tag | vSphere Tag | This property displays the vSphere Tag Name. |
| summary\|parentCluster | Parent Cluster | This property displays the Parent Cluster. |
| summary\|parentDatacenter | Parent Datacenter | This property displays the Parent Datacenter. |
| summary\|parentVcenter | Parent Vcenter | This property displays the Parent Vcenter. |

**Table 555: Datastore Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| datastore\|maxObservedNumberRead | Datastore I/O\|Highest Observed Number of Read Requests | |
| datastore\|maxObservedNumberWrite | Datastore I/O\|Highest Observed Number of Write Requests | |
| datastore\|maxObservedOIO | Datastore I/O\|Highest Observed Outstanding Requests | |
| datastore\|maxObservedRead | Datastore I/O\|Highest Observed Read Rate(kbps) | |
| datastore\|maxObservedWrite | Datastore I/O\|Highest Observed Write Rate(kbps) | |
| net\|discoveryProtocol\|cdp\|timeToLive | Network I/O\|Discovery Protocol\|Cisco Discovery Protocol\|Time to Live | |
| net\|discoveryProtocol\|lldp\|timeToLive | Network I/O\|Discovery Protocol\|Link Layer Discovery Protocol\|Time to Live | |

Datastore properties collected for host system objects have been disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

**Table 556: Storage Path Properties Collected for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| storageAdapter\|port_WWN | Storage Adapter\|Port WWN | The port world wide name for storage adapter. Available for FC adapters only. |

**Table 557: Compliance Configuration Related Properties for Host System Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|security\|password_max_days | Configuration\|Security\| Password Max Days | Password Max Days |
| config\|security\|welcome_message | Configuration\|Security\| Welcome Message Configured | Welcome Message Configured |
| config\|security\|issue | Configuration\|Security\|SSH Connection Banner Message Configured | SSH Connection Banner Message Configured |
| config\|security\| host_client_session_timeout | Configuration\|Security\|Host Client Session Timeout | Host Client Session Timeout |
| config\|security\| has_lockdown_exception_users | Configuration\|Security\|Has Lockdown exception users | Has Lockdown exception users |

### Cluster Compute Resource Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration and summary properties for cluster compute resource objects.

**Table 558: License Properties for Cluster Objects**

| Property Key | Property Name | Description |
|---|---|---|
| License type | License Type | Displays the license type for the cluster object |
| Expiry | Expiry | Displays the number of days remaining before the license expires the cluster object<br>**NOTE**<br>An alert is generated if the license threshold is<br>• >= 80% - Warning<br>• >= 90% - Immediate<br>• =95% - Catastrophic |

**Table 559: Configuration Properties Collected for Cluster Compute Resource Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|name | Name | Name |

**Table 560: Summary Properties Collected for Cluster Compute Resource Objects**

| Property Key | Property Name | Description |
|---|---|---|
| Summary\|Is Horizon Managed | Is Horizon Managed | This property displays whether the selected object is managed by Horizon or not. |
| summary\|parentDatacenter | Parent data center | This property displays the Parent data center. |
| summary\|parentVcenter | Parent vCenter | This property displays the Parent vCenter. |
| summary\|customTag\|customTagValue | Value | This property displays the Custom Tag Value. |
| summary\|tag | vSphere Tag | This property displays the vSphere Tag Name. |

**Table 561: DR, DAS, and DPM Configuration Properties Collected for Cluster Compute Resource Objects**

| Property Key | Property Name | Description |
|---|---|---|
| configuration\|drsconfig\|enabled | Enabled | Indicates whether DRS is enabled |
| configuration\|drsconfig\|<br>defaultVmBehavior | Default DRS Behavior | Default DRS Behavior |
| configuration\|drsconfig\|affinityRules | Affinity Rules | DRS Affinity Rules |
| configuration\|dasconfig\|enabled | HA Enabled | HA Enabled |
| configuration\|dasconfig\|<br>admissionControlEnabled | Admission Control Enabled | Admission Control Enabled |
| configuration\|dpmconfiginfo\|enabled | DPM Enabled | DPM Enabled |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| configuration\|dpmconfiginfo\| defaultDpmBehavior | Default DPM Behavior | Default DPM Behavior |
| configuration\|infraUpdateHaConfig\| remediation | Cluster Configuration\|HA Configuration\|Remediation | This property displays the Remediation mode taken by vSphere Cluster to deal with host failure. The options available are Quarantine mode, Maintenance mode, Mixed mode. |
| configuration\|drsConfig\| pctIdleMBInMemDemand | Cluster Configuration\|DRS Configuration\|Idle Consumed Memory | |
| configuration\|drsConfig\|targetBalance | Cluster Configuration\|DRS Configuration\|Tolerable imbalance threshold | |

DRS properties are collected for disaster recovery. DAS properties are collected for high availability service, formerly distributed availability service. DPM properties are collected for distributed power management.

### Resource Pool Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, CPU, memory, and summary properties for resource pool objects.

**Table 562: Configuration Properties Collected for Resource Pool Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|name | Name | Name |
| config\|cpuAllocation\|reservation | Reservation | CPU reservation |
| config\|cpuAllocation\|limit | Limit | CPU limit |
| config\|cpuAllocation\| expandableReservation | Expandable Reservation | CPU expandable reservation |
| config\|cpuAllocation\|shares\|shares | Shares | CPU shares |
| config\|memoryAllocation\|reservation | Reservation | Memory reservation |
| config\|memoryAllocation\|limit | Limit | Memory limit |
| config\|memoryAllocation\| expandableReservation | Expandable Reservation | Memory expandable reservation |
| config\|memoryAllocation\|shares\|shares | Shares | Memory shares |

**Table 563: CPU Usage Properties Collected for Resource Pool Objects**

| Property Key | Property Name | Description |
|---|---|---|
| cpu\|limit | CPU Limit | CPU Limit |
| cpu\|reservation | CPU reservation | CPU Reservation |
| cpu\|expandable_reservation | CPU expandable reservation | CPU Expandable Reservation |
| cpu\|shares | CPU Shares | CPU Shares |
| cpu\|corecount_provisioned | Provisioned vCPU(s) | Number of CPUs. It counts both the vSocket and vCore. A VM with 2 |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| | | vSockets x 4 vCores each has 8 vCPU. |

**Table 564: Memory Properties Collected for Resource Pool Objects**

| Property Key | Property Name | Description |
|---|---|---|
| mem\|limit | Memory limit | Memory limit |
| mem\|reservation | Memory reservation | Memory reservation |
| mem\|expandable_reservation | Memory expandable reservation | Memory expandable reservation |
| mem\|shares | Memory Shares | Memory Shares |

**Table 565: Summary Properties Collected for Resource Pool Objects**

| Property Key | Property Name | Description |
|---|---|---|
| summary\|customTag\|customTagValue | Value | Custom Tag Value |
| summary\|tag | vSphere Tag | vSphere Tag Name |

### Data Center Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration and summary properties for data center objects.

**Table 566: Configuration Properties Collected for Data Center Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|name | Name | Name |

**Table 567: Summary Properties Collected for Data Center Objects**

| Property Key | Property Name | Description |
|---|---|---|
| summary\|parentVcenter | Parent Vcenter | Parent Vcenter |
| summary\|customTag\|customTagValue | Value | Custom Tag Value |
| summary\|tag | vSphere Tag | vSphere Tag Name |

### Storage Pod Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration and summary properties for storage pod objects.

**Table 568: Configuration Properties Collected for Storage Pod Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|name | Name | Name |
| config\|sdrsconfig\| vmStorageAntiAffinityRules | VM storage antiaffinity rules | Storage Distributed Resource Scheduler (SDRS) VM anti-affinity rules |
| config\|sdrsconfig\|vmdkAntiAffinityRules | VMDK antiaffinity rules | Storage Distributed Resource Scheduler (SDRS) Virtual Machine Disk (VMDK) anti-affinity rules |

### VMware Distributed Virtual Switch Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration and summary properties for VMware distributed virtual switch objects.

**Table 569: Configuration Properties Collected for VMware Distributed Virtual Switch Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\| networkResourceManagementEnabled | Configuration\|Network IO Control | This property shows the status of Network IO Control. Enabled means control over Network IO (using shares, limitations and reservation) are in place on each port group. |
| config\|name | Name | Name |

**Table 570: Capability Properties Collected for VMware Distributed Virtual Switch Objects**

| Property Key | Property Name | Description |
|---|---|---|
| capability\|nicTeamingPolicy | NIC Teaming Policy | NIC Teaming Policy |

### Distributed Virtual Port Group Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration and summary properties for distributed virtual port group objects.

**Table 571: Configuration Properties Collected for Distributed Virtual Port Group Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|portBinding | Configuration\|Port Binding | This property shows how ports are assigned to virtual machines connected to this distributed port group. Available values: earlyBinding, ephemeral, lateBinding. |
| config\|portAllocation | Configuration\|Port Allocation | This property shows the type of port allocation, such as elastic or fixed. Options are: true=Elastic, false=Fixed. |
| config\|name | Name | Name |
| Configuration\|Uplink | Uplink | Indicates whether the portgroup is uplink portgroup. |

**Table 572: Summary Properties Collected for Distributed Virtual Port Group Objects**

| Property Key | Property Name | Description |
|---|---|---|
| summary\|active_uplink_ports | Active DV uplinks | Active DV uplinks |

### Datastore Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration, summary, and properties about datastore use for datastore objects.

**Table 573: Configuration Properties of Datastore and Datastore Cluster Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|iormConfigStatus | Configuration\|Storage IO Control Status | Displays the status of Storage IO Control. If enabled control over Disk IO (using shares, limitations and reservation) is in place.Value: True or False. |
| summary\|total_number_datastores | Summary\|Total Number of Datastores | Displays the total number of member datastores in the cluster. |
| summary\|parentVcenter | Summary\|Parent vCenter | Displays the details of the parent vCenter. |
| summary\|parentDatacenter (GB) | Summary\|Parent Datacenter | Displays the details of the parent Datacenter. |

**Table 574: Capacity Properties Collected for vSAN Datastore Objects**

| Property Key | Property Name | Description |
|---|---|---|
| Capacity\|Available Space (GB) | Available Space | Displays the available disk space in GB. |
| Capacity\|Provisioned (GB) | Provisioned (GB) | Displays the provisioned datastore size in GB. |
| Capacity\|Total Capacity (GB) | Total Capacity (GB) | Displays the total datastore capacity in GB. |
| Capacity\|Total Provisioned Consumer Space (GB) | Total Provisioned Consumer Space (GB) | Displays the total provisioned consumer space in GB. |
| Capacity\|Used Space (GB) | Used Space (GB) | Displays the used disk space in GB. |
| Capacity\|Used Space (%) | Used Space (%) | Displays the used disk space in percentage. |
| Capacity\|Usable Capacity (GB) | Usable Capacity (GB) | Displays the usable disk capacity in GB. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| | | **NOTE**<br>Earlier the vSAN Datastore base rate was calculated on the basis of Total Capacity of the disk, now the vSAN datastore base rate is calculated based on the usable capacity. |

**Table 575: Summary Properties Collected for Datastore Objects**

| Property Key | Property Name | Description |
|---|---|---|
| summary\|vmfs_version | VMFS (Virtual Machine File System) Version | Displays the VMFS version number, contains both major version and minor version number.<br>**NOTE**<br>The VMFS version property is visible, only when the datastore type is VMFS. |
| summary\|diskCapacity | Disk Capacity | Disk Capacity |
| summary\|isLocal | Is Local | Is local datastore |
| summary\|customTag\|customTagValue | Value | Custom Tag Value |
| summary\|accessible | Datastore Accessible | Datastore Accessible |
| summary\|path | Summary\|Path | |
| summary\|scsiAdapterType | Summary\|SCSI Adapter Type | This property is disabled by default. |
| summary\|aliasOf | Summary\|Alias Of | Indicates whether the datastore is an alias of another. The published value is the container ID of the datastore for which it is an alias.<br>**NOTE**<br>This property may have 2 values. It's either "none", that means the datastore is not an alias of another datastore, or datastore <containerID> that is the Container ID of the datastore for which this is an alias. |

**Table 576: Datastore Properties Collected for Datastore Objects**

| Property Key | Property Name | Description |
|---|---|---|
| datastore\|hostcount | Host Count | Host Count |
| datastore\|hostScsiDiskPartition | Host SCSI Disk Partition | Host SCSI Disk Partition |
| * datastore\|<br>maxObservedNumberRead | Datastore I/O\|Highest Observed Number of Read Requests | Disabled |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| * datastore\| maxObservedNumberWrite | Datastore I/O\|Highest Observed Number of Write Requests | Disabled |
| * datastore\|maxObservedOIO | Datastore I/O\|Highest Observed Outstanding Requests | Disabled |
| * datastore\|maxObservedRead | Datastore I/O\|Highest Observed Read Latency | Disabled |
| * datastore\| maxObservedReadLatency | Datastore I/O\|Highest Observed Read Latency | Disabled |
| * datastore\|maxObservedWrite | Datastore I/O\|Highest Observed Write Latency | Disabled |
| * datastore\| maxObservedWriteLatency | Datastore I/O\|Highest Observed Write Latency | Disabled |

**Table 577: Datastore Properties Collected for vVol Datastore Objects**

| Property Key | Property Name | Description |
|---|---|---|
| storageArray\|modelId | Storage Array\|Model | Storage array model of VVol datastore.<br>**NOTE**<br>This property is published for vVol datastores only and is available starting from vCenter version 6.0. |
| storageArray\|name | Storage Array\|Name | Storage array name of vVol datastore.<br>**NOTE**<br>This property is published for vVol datastores only and is available starting from vCenter version 6.0. |
| storageArray\|id | Storage Array\|ID | Storage array ID of vVol datastore.<br>**NOTE**<br>This property is published for vVol datastores only and is available starting from vCenter version 6.0. |
| storageArray\|vendorId | Storage Array\|Vendor | Storage array vendor of vVol datastore.<br>**NOTE**<br>This property is published for vVol datastores only and is available starting from vCenter version 6.0. |
| protocolEndpoints\|name | Protocol Endpoints\|Name | Protocol endpoint's name of vVol datastore. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| | | **NOTE**<br>This is an instanced property that is published per protocol endpoint instance (e. g. eui.3362663138636633) for vVol datastores only. It is available starting from vCenter version 6.0. |
| protocolEndpoints\|type | Protocol Endpoints\|Type | Protocol endpoint's type of vVol datastore.<br>**NOTE**<br>This is an instanced property that is published per protocol endpoint instance (e. g. eui.3362663138636633) for vVol datastores only. It is available starting from vCenter version 6.5. |
| protocolEndpoints\|hosts | Protocol Endpoints\|Hosts | Hosts associated with protocol endpoint of vVol datasore.<br>**NOTE**<br>This is an instanced property that is published per protocol endpoint instance (e. g. eui.3362663138636633) for vVol datastores only. It is available starting from vCenter version 6.0. |

Datastore properties marked with an asterisk (*) have been disabled in this version of VMware Aria OperationsVMware Cloud Foundation Operations. This means that they do not collect data by default.

### vSphere Pod Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects summary and event properties for vSphere Pods.

**Table 578: Summary Properties Collected for vSphere Pod Objects**

| Property Key | Localized Name | Description |
|---|---|---|
| config\|name | Configuration\|Name | Resource name. |
| config\|guestFullName | Configuration\|Guest OS from vCenter | This is the value provided by vCenter. vCenter set it during VM creation. The value may not match the value inside the Guest. |
| config\|version | Configuration\|Version | Virtual Machine Version. |
| config\|createDate | Configuration\|Creation Date | Object Creation Date. |
| config\|numVMDKs | Configuration\|Number of Virtual Disks | Number of Virtual Disks. |
| config\|faultTolerant | Configuration\|Fault Tolerant | Fault tolerance enabled. |
| config\|ft_role | Configuration\|FT Role | Role of the VM in Fault Tolerance Group. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| config\|ft_peer_vm | Configuration\|FT Peer VM | Peer of the VM in Fault Tolerance Group. |
| config\|hardware\|numCpu | Configuration\|Hardware\| Number of virtual CPUs | Number of virtual CPUs. |
| config\|hardware\|memoryKB | Configuration\|Hardware\| Memory | Memory. |
| config\|hardware\|thinEnabled | Configuration\|Hardware\|Thin Provisioned Disk | Thin Provisioned Disk. |
| config\|hardware\|numCoresPerSocket | Configuration\|Hardware\| Number of CPU cores per socket | Number of CPU cores per virtual socket. |
| config\|hardware\|numSockets | Configuration\|Hardware\| Number of virtual sockets | Number of virtual sockets. |
| config\|hardware\|diskSpace | Configuration\|Hardware\|Disk Space | Disk space metrics. |
| config\|cpuAllocation\|reservation | Configuration\|CPU Resource Allocation\|Reservation | N/A |
| config\|cpuAllocation\|limit | Configuration\|CPU Resource Allocation\|Limit | |
| config\|cpuAllocation\|shares\|shares | Configuration\|CPU Resource Allocation\|Shares\|Shares | |
| config\|memoryAllocation\|reservation | Configuration\|Memory Resource Allocation\| Reservation | |
| config\|memoryAllocation\|limit | Configuration\|Memory Resource Allocation\|Limit | |
| config\|memoryAllocation\|shares\|shares | Configuration\|Memory Resource Allocation\|Shares\| Shares | |
| config\|extraConfig\|mem_hotadd | Configuration\|Extra Configuration\|Memory Hot Add | Memory Hot Add Configuration. |
| config\|extraConfig\|vcpu_hotadd | Configuration\|Extra Configuration\|vCPU Hot Add | vCPU Hot Add Configuration. |
| config\|extraConfig\|vcpu_hotremove | Configuration\|Extra Configuration\|vCPU Hot Remove | vCPU Hot Remove Configuration. |
| config\|extraConfig\|mem_tps_share | Configuration\|Extra Configuration\|VM MEM TPS | N/A |
| config\|security\|disable_autoinstall | Configuration\|Security\| Disable tools auto install (isolation.tools.autoInstall.disable) | |
| config\|security\|disable_console_copy | Configuration\|Security\| Disable console copy operations (isolation.tools.copy.disable) | |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| config\|security\|disable_console_dnd | Configuration\|Security\| Disable console drag and drop operations (isolation.tools.dnd.disable) | |
| config\|security\| enable_console_gui_options | Configuration\|Security\|Enable console GUI operations (isolation.tools.setGUIOptions .enable) | |
| config\|security\|disable_console_paste | Configuration\|Security\| Disable console paste operations (isolation.tools.paste.disable) | |
| config\|security\| disable_disk_shrinking_shrink | Configuration\|Security\| Disable virtual disk shrink (isolation.tools.diskShrink.dis able) | |
| config\|security\| disable_disk_shrinking_wiper | Configuration\|Security\| Disable virtual disk wiper (isolation.tools.diskWiper.disa ble) | |
| config\|security\|disable_hgfs | Configuration\|Security\| Disable HGFS file transfers (isolation.tools.hgfsServerSet. disable) | |
| config\|security\| disable_independent_nonpersistent | Configuration\|Security\|Avoid using independent nonpersistent disks (scsiX:Y.mode) | |
| config\|security\|enable_intervm_vmci | Configuration\|Security\|Enable VM-to-VM communication through VMCI (vmci0.unrestricted) | |
| config\|security\|enable_logging | Configuration\|Security\|Enable VM logging (logging) | |
| config\|security\|disable_monitor_control | Configuration\|Security\| Disable VM Monitor Control (isolation.monitor.control.disa ble) | |
| config\|security\| enable_non_essential_3D_features | Configuration\|Security\|Enable 3D features on Server and desktop virtual machines (mks.enable3d) | |
| config\|security\| disable_unexposed_features_autologon | Configuration\|Security\| Disable unexposed features - autologon (isolation.tools.ghi.autologon. disable) | |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| config\|security\|disable_unexposed_features_biosbbs | Configuration\|Security\|Disable unexposed features - biosbbs (isolation.bios.bbs.disable) | |
| config\|security\|disable_unexposed_features_getcreds | Configuration\|Security\|Disable unexposed features - getcreds (isolation.tools.getCreds.disable) | |
| config\|security\|disable_unexposed_features_launchmenu | Configuration\|Security\|Disable unexposed features - launchmenu (isolation.tools.ghi.launchmenu.change) | |
| config\|security\|disable_unexposed_features_memsfss | Configuration\|Security\|Disable unexposed features - memsfss (isolation.tools.memSchedFakeSampleStats.disable) | |
| config\|security\|disable_unexposed_features_protocolhandler | Configuration\|Security\|Disable unexposed features - protocolhandler (isolation.tools.ghi.protocolhandler.info.disable) | |
| config\|security\|disable_unexposed_features_shellaction | Configuration\|Security\|Disable unexposed features - shellaction (isolation.ghi.host.shellAction.disable) | |
| config\|security\|disable_unexposed_features_toporequest | Configuration\|Security\|Disable unexposed features - toporequest (isolation.tools.dispTopoRequest.disable) | |
| config\|security\|disable_unexposed_features_trashfolderstate | Configuration\|Security\|Disable unexposed features - trashfolderstate (isolation.tools.trashFolderState.disable) | |
| config\|security\|disable_unexposed_features_trayicon | Configuration\|Security\|Disable unexposed features - trayicon (isolation.tools.ghi.trayicon.disable) | |
| config\|security\|disable_unexposed_features_unity | Configuration\|Security\|Disable unexposed features - unity (isolation.tools.unity.disable) | |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| config\|security\|disable_unexposed_features_unity_interlock | Configuration\|Security\|Disable unexposed features - unity-interlock (isolation.tools.unityInterlockOperation.disable) | |
| config\|security\|disable_unexposed_features_unity_taskbar | Configuration\|Security\|Disable unexposed features - unity-taskbar (isolation.tools.unity.taskbar.disable) | |
| config\|security\|disable_unexposed_features_unity_unityactive | Configuration\|Security\|Disable unexposed features - unity-unityactive (isolation.tools.unityActive.disable) | |
| config\|security\|disable_unexposed_features_unity_windowcontents | Configuration\|Security\|Disable unexposed features - unity-windowcontents (isolation.tools.unity.windowContents.disable) | |
| config\|security\|disable_unexposed_features_unitypush | Configuration\|Security\|Disable unexposed features - unitypush (isolation.tools.unity.push.update.disable) | |
| config\|security\|disable_unexposed_features_versionget | Configuration\|Security\|Disable unexposed features - versionget (isolation.tools.vmxDnDVersionGet.disable) | |
| config\|security\|disable_unexposed_features_versionset | Configuration\|Security\|Disable unexposed features - versionset (solation.tools.guestDnDVersionSet.disable) | |
| config\|security\|disable_vix_messages | Configuration\|Security\|Disable VIX messages from the VM (isolation.tools.vixMessage.disable) | |
| config\|security\|enable_vga_only_mode | Configuration\|Security\|Disable all but VGA mode on virtual machines (svga.vgaOnly) | |
| config\|security\|limit_console_connection | Configuration\|Security\|Limit number of console connections (RemoteDisplay.maxConnection) | |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| config\|security\|limit_log_number | Configuration\|Security\|Limit number of log files (log.keepOld) | |
| config\|security\|limit_log_size | Configuration\|Security\|Limit log file size (log.rotateSize) | |
| config\|security\|limit_setinfo_size | Configuration\|Security\|Limit VMX file size (tools.setInfo.sizeLimit) | |
| config\|security\|enable_console_VNC | Configuration\|Security\|Enable access to VM console via VNC protocol (RemoteDisplay.vnc.enabled) | |
| config\|security\| disable_device_interaction_connect | Configuration\|Security\| Disable unauthorized removal, connection of devices (isolation.device.connectable.disable) | |
| config\|security\| disable_device_interaction_edit | Configuration\|Security\| Disable unauthorized modification of devices (isolation.device.edit.disable) | |
| config\|security\|enable_host_info | Configuration\|Security\|Enable send host information to guests (tools.guestlib.enableHostInfo) | |
| config\|security\|network_filter_enable | Configuration\|Security\|Enable dvfilter network APIs (ethernetX.filterY.name) | |
| config\|security\| vmsafe_cpumem_agentaddress | Configuration\|Security\| VMsafe CPU/memory APIs - IP address (vmsafe.agentAddress) | |
| config\|security\| vmsafe_cpumem_agentport | Configuration\|Security\| VMsafe CPU/memory APIs - port number (vmsafe.agentPort) | |
| config\|security\|vmsafe_cpumem_enable | Configuration\|Security\|Enable VMsafe CPU/memory APIs (vmsafe.enable) | |
| config\|security\| disconnect_devices_floppy | Configuration\|Security\| Disconnect floppy drive | |
| config\|security\|disconnect_devices_cd | Configuration\|Security\| Disconnect CD-ROM | |
| config\|security\|disconnect_devices_usb | Configuration\|Security\| Disconnect USB controller | |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| config\|security\| disconnect_devices_parallel | Configuration\|Security\| Disconnect parallel port | |
| config\|security\| disconnect_devices_serial | Configuration\|Security\| Disconnect serial port | |
| config\|security\|pci_device_configured | Configuration\|Security\|DCUI timeout | |
| runtime\|memoryCap | Runtime\|Memory Capacity | Memory Capacity. |
| cpu\|limit | CPU\|CPU Limit | CPU Limit. |
| cpu\|reservation | CPU\|CPU reservation | CPU Reservation. |
| cpu\|speed | CPU\|CPU | CPU Speed. |
| mem\|host_reservation | Memory\|Host Active | Machine Active. |
| mem\|host_active | Memory\|Host Usage | Machine Usage. |
| net\|mac_address | Network\|Mac Address | N/A |
| net\|ip_address | Network\|IP Address | |
| net\|subnet_mask | Network\|Subnet Mask | |
| net\|ipv6_address | Network\|IPv6 Address | IPv6 Address. |
| net\|ipv6_prefix_length | Network\|IPv6 Prefix Length | IPv6 Prefix Length. |
| net\|default_gateway | Network\|Default Gateway | N/A |
| net\|nvp_vm_uuid | Network\|NVP VM UUID | |
| net\|vnic_type | Network\|Virtual NIC Type | Virtual Machine's network adapter type. |
| net\|vnic_label | Network\|Label | Device label. |
| summary\|UUID | Summary\|UUID | Instance UUID in vCenter that uniquely identify all virtual machine instances. |
| summary\|MOID | Summary\|MOID | Managed object ID in vCenter. This is unique in scope of vCenter. |
| summary\|swapOnlyDatastore | Summary\|Datastore with only swap file | Datastore containing only the swap file and no other files from this VM. |
| summary\|customTag\|customTagValue | Summary\|Custom Tag\|Value | Custom Tag Value. |
| summary\|tag | Summary\|vSphere Tag | vSphere Tag Name. |
| summary\|tagJson | Summary\|vSphere Tag Json | vSphere Tag in Json format. |
| summary\|folder | Summary\|vSphere Folder | vSphere Folder Name. |
| summary\|parentCluster | Summary\|Parent Cluster | Parent Cluster. |
| summary\|parentHost | Summary\|Parent Host | Parent Host. |
| summary\|parentDatacenter | Summary\|Parent Datacenter | Parent Datacenter. |
| summary\|parentNamespace | Summary\|Parent Namespace | Parent Namespace. |
| summary\|parentVcenter | Summary\|Parent vCenter | Parent vCenter. |
| summary\|parentFolder | Summary\|Parent Folder | Parent Folder. |
| summary\|datastore | Summary\|Datastore(s) | Datastore(s). |
| summary\|guest\|fullName | Summary\|Guest Operating System\|Guest OS from Tools | This is the value provided by VMware Tools. This value will differ to the value set in vCenter if the Guest OS was upgraded, or a different Guest OS was installed. |
| summary\|guest\|ipAddress | Summary\|Guest Operating System\|Guest OS IP Address | Guest OS IP Address. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| summary\|guest\|hostName | Summary\|Guest Operating System\|Hostname | Hostname of the guest operating system, if known. |
| summary\|guest\|toolsRunningStatus | Summary\|Guest Operating System\|Tools Running Status | Guest Tools Running Status. |
| summary\|guest\|toolsVersionStatus2 | Summary\|Guest Operating System\|Tools Version Status | Guest Tools Version Status 2. |
| summary\|guest\|toolsVersion | Summary\|Guest Operating System\|Tools Version | VM tools version installed on guest OS. |
| summary\|guest\| vrealize_operations_agent_id | Summary\|Guest Operating System\|vRealize Operations Agent ID | An ID to identify a VM in Agent Adapter's world. |
| summary\|guest\| vrealize_operations_euc_agent_id | Summary\|Guest Operating System\|vRealize Operations Euc Agent ID | An ID to identify a VM in Agent Adapter's world. |
| summary\|config\|numEthernetCards | Summary\|Configuration\| Number of NICs | Number of NICs. |
| summary\|config\|productName | Summary\|Configuration\| Product Name | Product Name. |
| summary\|config\|appliance | Summary\|Configuration\| Appliance | Appliance. |
| summary\|runtime\|isIdle | Summary\|Runtime\|Idleness indicator | This property indicates whether the monitored instance is idle or not. |
| summary\|runtime\|powerState | Summary\|Runtime\|Power State | Power State. |
| summary\|runtime\|connectionState | Summary\|Runtime\| Connection State | Connection State. |
| summary\|smbiosUUID | SMBIOS UUID | System Management BIOS UUID of a virtual machine.<br>**NOTE**<br>The SMBIOS UUID metric for vSphere Pod is disabled by default. You have to enable the metric at the policy level. |
| guestfilesystem\|capacity_property | Guest File System\|Guest File System Capacity Property | Total capacity of guest file system as a property. |
| guestfilesystem\|capacity_property_total | Guest File System\|Total Capacity Property | Total capacity of guest file system as a property. |
| virtualDisk\|datastore | Virtual Disk\|Datastore | Datastore. |
| virtualDisk\|configuredGB | Virtual Disk\|Configured | Virtual Disk configured disk space. |
| virtualDisk\|label | Virtual Disk\|Label | Device Label. |
| virtualDisk\|fileName | Virtual Disk\|File Name | Virtual Disk file name. |
| diskspace\|snapshot\|mor | Disk Space\|Snapshot\| Managed Object Reference | Managed Object Reference. |
| diskspace\|snapshot\|name | Disk Space\|Snapshot\|Name | Snapshot name. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| diskspace\|snapshot\|numberOfDays | Disk Space\|Snapshot\|Number of Days Old | Number of days since snapshot creation. |
| diskspace\|snapshot\|snapshotAge | Disk Space\|Snapshot\|Age (Days) | Virtual Machine's topmost snapshot age in days. |
| diskspace\|snapshot\|creator | Disk Space\|Snapshot\|Creator | Creator. |
| diskspace\|snapshot\|description | Disk Space\|Snapshot\| Description | Snapshot description. |
| vsan\|policy\|compliance | vSAN\|VM Storage Policies\| Compliance | Compliance status of the VM storage object. |
| datastore\|maxObservedNumberRead | Datastore\|Highest Observed Number of Read Requests | Highest Observed Number of Read Requests. |
| datastore\|maxObservedRead | Datastore\|Highest Observed Read Rate | Highest Observed Read Rate (KBps). |
| datastore\|maxObservedNumberWrite | Datastore\|Highest Observed Number of Write Requests | Highest Observed Number of Write Requests. |
| datastore\|maxObservedWrite | Datastore\|Highest Observed Write Rate | Highest Observed Write Rate (KBps). |
| datastore\|maxObservedOIO | Datastore\|Highest Observed Outstanding Requests | Highest Observed Outstanding Requests. |

**Table 579: Compliance Configuration Related Properties for vSphere Pod Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|security\| disconnect_devices_virtualhdaudiocard | Configuration\|Security\|Virtual HD Audio Card Disconnected | NIL |
| config\|security\| disconnect_devices_virtualahcicontroller | Configuration\|Security\|Virtual AHCI Controller Disconnected | NIL |
| config\|security\| disconnect_devices_virtualensoniq1371 | Configuration\|Security\|Virtual Ensoniq 1371 Disconnected | NIL |

### Namespace Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects summary and event properties for Namespace.

**Table 580: Summary Properties Collected for Namespace Objects**

| Property Key | Localized Name | Description |
|---|---|---|
| config\|name | Configuration\|Name | Resource name |
| config\|resourceLimits\|namespace\|cpu | Configuration\|Resource Limits\|Namespaces\|CPU | CPU |
| config\|resourceLimits\|namespace\|mem | Configuration\|Resource Limits\|Namespaces\|Memory | Memory |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| config\|resourceLimits\|namespace\|diskspace | Configuration\|Resource Limits\|Namespaces\|Disk Space | Disk space metrics |
| config\|resourceLimits\|containers\|cpu_request | Configuration\|Resource Limits\|Containers\|CPU Request | CPU Request Default |
| config\|resourceLimits\|containers\|cpu_limit | Configuration\|Resource Limits\|Containers\|CPU Limit | CPU Limit Default |
| config\|resourceLimits\|containers\|mem_request | Configuration\|Resource Limits\|Containers\|Memory Request | Memory Request Default |
| config\|resourceLimits\|containers\|mem_limit | Configuration\|Resource Limits\|Containers\|Memory Limit | Memory Limit Default |
| config\|objectLimits\|compute\|pod_count | Configuration\|Object Limits\|Compute\|Pods | Number of Pods |
| config\|objectLimits\|compute\|deployment_count | Configuration\|Object Limits\|Compute\|Deployments | Deployments |
| config\|objectLimits\|compute\|job_count | Configuration\|Object Limits\|Compute\|Jobs | Jobs |
| config\|objectLimits\|compute\|daemon_sets | Configuration\|Object Limits\|Compute\|Daemon Sets | Daemon Sets |
| config\|objectLimits\|compute\|replica_sets | Configuration\|Object Limits\|Compute\|Replica Sets | Replica Sets |
| config\|objectLimits\|compute\|replication_controllers | Configuration\|Object Limits\|Compute\|Replication Controllers | Replication Controllers |
| config\|objectLimits\|compute\|stateful_sets | Configuration\|Object Limits\|Compute\|Stateful Sets | Stateful Sets |
| config\|objectLimits\|storage\|config_maps | Configuration\|Object Limits\|Storage\|Config Maps | Config Maps |
| config\|objectLimits\|storage\|secret_count | Configuration\|Object Limits\|Storage\|Secrets | Secrets |
| config\|objectLimits\|storage\|persistent_volume_claim | Configuration\|Object Limits\|Storage\|Persistent Volume Claim | Persistent Volume Claim |
| config\|objectLimits\|network\|services | Configuration\|Object Limits\|Network\|Services | Services |
| summary\|parentDatacenter | Summary\|Parent Datacenter | Parent Datacenter |
| summary\|parentCluster | Summary\|Parent Cluster | Parent Cluster |
| summary\|parentVcenter | Summary\|Parent vCenter | Parent vCenter |
| mem\|limit | Memory\|Memory limit | Memory limit |
| mem\|reservation | Memory\|Memory reservation | Memory reservation |
| mem\|expandable_reservation | Memory\|Memory expandable reservation | Memory Expandable Reservation |
| mem\|shares | Memory\|Memory Shares | Memory Shares |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| cpu\|limit | CPU\|CPU Limit | CPU Limit |
| cpu\|reservation | CPU\|CPU Reservation | CPU Reservation |
| cpu\|expandable_reservation | CPU\|CPU expandable reservation | CPU expandable Reservation |
| cpu\|shares | CPU\|CPU Shares | CPU Shares |
| cpu\|corecount_provisioned | CPU\|Provisioned vCPU(s) | Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU. |

### Tanzu Kubernetes cluster Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects summary and event properties forTanzu Kubernetes clusters.

**Table 581: Summary Properties Collected for Tanzu Kubernetes cluster Objects**

| Property Key | Localized Name | Description |
|---|---|---|
| config\|name | Configuration\|Name | Resource name |
| config\|cpuAllocation\|reservation | Configuration\|CPU Resource Allocation\|Reservation | N/A |
| config\|cpuAllocation\|limit | Configuration\|CPU Resource Allocation\|Limit | N/A |
| config\|cpuAllocation\| expandableReservation | Configuration\|CPU Resource Allocation\|Expandable Reservation | N/A |
| config\|cpuAllocation\|shares\|shares | Configuration\|CPU Resource Allocation\|Shares\|Shares | N/A |
| config\|memoryAllocation\|reservation | Configuration\|Memory Resource Allocation\| Reservation | N/A |
| config\|memoryAllocation\|limit | Configuration\|Memory Resource Allocation\|Limit | N/A |
| config\|memoryAllocation\| expandableReservation | Configuration\|Memory Resource Allocation\| Expandable Reservation | N/A |
| config\|memoryAllocation\|shares\|shares | Configuration\|Memory Resource Allocation\|Shares\| Shares | N/A |
| cpu\|limit | CPU\|CPU Limit | CPU Limit |
| cpu\|reservation | CPU\|CPU Reservation | CPU Reservation |
| cpu\|expandable_reservation | CPU\|CPU expandable reservation | CPU expandable Reservation |
| cpu\|shares | CPU\|CPU Shares | CPU Shares |

*Table continued on next page*

*Continued from previous page*

| Property Key | Localized Name | Description |
|---|---|---|
| cpu\|corecount_provisioned | CPU\|Provisioned vCPU(s) | Number of CPUs. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU. |
| mem\|limit | Memory\|Memory limit | Memory limit |
| mem\|reservation | Memory\|Memory reservation | Memory reservation |
| mem\|expandable_reservation | Memory\|Memory expandable reservation | Memory Expandable Reservation |
| mem\|shares | Memory\|Memory Shares | Memory Shares |
| summary\|parentDatacenter | Summary\|Parent Datacenter | Parent Datacenter |
| summary\|parentNamespace | Summary\|Parent Namespace | Parent Namespace |

### All Folder Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects configuration and summary properties for All Folders.

**Table 582: Summary Properties Collected for All Folder Objects**

| Property Key | Property Name | Description |
|---|---|---|
| summary\|parentDatacenter | Summary\|Parent Datacenter | This property shows the details of the parent datacenter. |
| summary\|parentVcenter | Summary\|Parent vCenter | This property shows the details of the parent vCenter. |
| summary\|tag | Summary\|vSphere Tag | This property shows the details of the vSphere tag name. |

### Self-Monitoring Properties for VMware Aria OperationsVMware Cloud Foundation Operations

VMware Aria OperationsVMware Cloud Foundation Operations uses the VMware Aria OperationsVMware Cloud Foundation Operations adapter to collect properties that monitor its own objects. These self-monitoring properties are useful for monitoring changes within VMware Aria OperationsVMware Cloud Foundation Operations.

### Analytics Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects properties for the VMware Aria OperationsVMware Cloud Foundation Operations analytics service.

**Table 583: Properties Collected for Analytics Service Objects**

| Property Key | Property Name | Description |
|---|---|---|
| HAEnabled | HA Enabled | Indicates HA is enabled with a value of 1, disabled with a value of 0. |

*Table continued on next page*

*Continued from previous page*

| Property Key | Property Name | Description |
|---|---|---|
| ControllerDBRole | Role | Indicates persistence service role for the controller: 0 – Primary, 1 – Replica, 4 – Client.. |
| ShardRedundancyLevel | Shard redundancy level | The target number of redundant copies for Object data. |
| LocatorCount | Locator Count | The number of configured locators in the system |
| ServersCount | Servers Count | The number of configured servers in the system |

### Node Properties

VMware Aria OperationsVMware Cloud Foundation Operations collects properties for the VMware Aria OperationsVMware Cloud Foundation Operations node objects.

**Table 584: Configuration Properties Collected for Node Objects**

| Property Key | Property Name | Description |
|---|---|---|
| config\|numCpu | Number of CPU | Number of CPUs |
| config\|numCoresPerCpu | Number of cores per CPU | Number of cores per CPU |
| config\|coreFrequency | Core Frequency | Core Frequency |

**Table 585: Memory Properties Collected for Node Objects**

| Property Key | Property Name | Description |
|---|---|---|
| mem\|RAM | System RAM | System RAM |

**Table 586: Service Properties Collected for Node Objects**

| Property Key | Property Name | Description |
|---|---|---|
| service\|proc\|pid | Process ID | Process ID |

### OS and Application Monitoring Properties

Properties are collected for operating systems, application services, remote checks, Linux processes, and Windows services which can be used to create reports, views, and dashboards.

### Guest Information Properties

VMware Aria OperationsVMware Cloud Foundation Operations displays the following guest information properties for all objects created by the OS and Application Monitoring management pack.

- Guest Info
  - Hostname
  - IP
  - OS Name

- OS Version
- Telegraf Version

Other properties of operating systems and application services are available under **Properties › Tags**.

### Service Discovery Properties

VMware Aria OperationsVMware Cloud Foundation Operations displays object properties for service discovery.

### Service Discovery Adapter Instance Properties

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the service discovery adapter instance.

**Table 587: Service Discovery Adapter Instance Properties**

| Property Name | Description |
|---|---|
| Action Identifier | An FQDN and IP pair of the end point vCenter Server that is used to identify the adapter instance that has to run actions on the vCenter. |

### Virtual Machine Properties

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for virtual machines.

**Table 588: Virtual Machine Properties**

| Property Name | Description |
|---|---|
| Guest OS Services\|Authentication Method | Refers to the VM guest operating system authentication method. The guest operating system can be authenticated either via a common user/password or a guest alias. |
| Guest OS Services\|Discovery Status | Reflects the result of service discovery operation on the VM's guest operating system. |
| Guest OS Services\|Authentication Status | Guest operating system authentication status. |
| Guest OS Services\|Inbound Ports | List of VM inbound ports. These are the ports on which the discovered services are listening. |
| SRM Info\|Protection Group | Protection group to which the VM belongs. |
| SRM Info\|Recovery Plans | List of recovery plans covering the VM. |

### Services Properties

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for services.

**Table 589: Services Properties**

| Property Name | Description |
|---|---|
| Type | The name of the service type. |
| Install Path | The install path. |
| Ports | List of service listening ports. |
| Virtual Machine | The name of the parent VM. |
| Virtual Machine MOID | The MOID of the VM. |
| Version | Version of the discovered service. |
| Is Application Member | Indicates that the service is a member of the group of services forming an application. |
| Category | Category of the service. |
| Connection Type | If there is a remote process that was connected to one of the listening ports of the given service, then the property's value is set to `Incoming`. If not, it is set to `Outgoing`. If there is no connection to another service, then the value of the property is set to `N/A`. |
| Has Dynamic Port | Indicates whether the service has dynamic ports or not. |
| Status | Indicates the status of the service.<br>**Up:** The service is running.<br><br>**Down:** The service is unavailable on the monitored VM.<br><br>**Unavailable:** The service is unavailable on a VM that is not being monitored.<br><br>**None:** The service is not available within 7 days. |

### Properties for vSAN

VMware Aria OperationsVMware Cloud Foundation Operations displays object properties for vSAN.

### Properties for vSAN Disk Groups

VMware Aria OperationsVMware Cloud Foundation Operations displays the following property for vSAN disk groups:

- vSAN Disk Groups: Configuration|vSAN Configuration
- vSAN Disk Groups: Configuration | Number of Disks

### Properties for vSAN Cluster

The VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for vSAN cluster.

| Property Name | Description |
|---|---|
| Configuration\|vSAN\|vSAN ESA | Indicates whether vSAN ESA configuration is enabled on the vSAN cluster. |
| Configuration\|vSAN\|Deduplication and Compression Enabled | Indicates whether deduplication and compression is eanbled on the vSAN cluster. |
| Configuration\|vSAN\|Preferred fault domain | Indicates whether the preferred fault domain is not set for the witness host in a vSAN Stretched cluster. |
| Configuration\|vSAN\|Stretched Cluster | Indicates whehter vSAN stretch cluster is enabled or not. |
| Configuration\|vSAN\|vSAN Configuration | Indicates whether the vSAN cluster is configured or not. |
| Configuration\|vSAN\|Encryption | Indicates whether the vSAN cluster is encypted or not. |
| Configuration \| vSAN \| File Service | Indicates whether vSAN File Services is enabled or not. |
| Configuration \| vSAN \| File Service Domain:<domainName> \| DNS Servers | Indicates the IP addresses of DNS servers, which are used to resolve the host names within the DNS domain. |
| Configuration \| vSAN \| File Service Domain:<domainName> \| DNS Suffixes | Indicates the list of DNS suffixes which can resolved by the DNS servers. |
| Configuration \| vSAN \| File Service Domain:<domainName> \| Gateway | Indicated the default gateway IP address for the file service access point. |
| Configuration \| vSAN \| File Service Domain:<domainName> \| Primary IP | Indicates the primary IP address for the file service. |
| Configuration \| vSAN \| File Service Domain:<domainName> \| Subnet Mask | Indicates the subnet mask for the vSAN cluster. |
| Summary \| Type | vSAN Cluster Type |
| Configuration \| vSAN \| File Service Domain:<domainName> \| IP Address :<ipaddress> \| FQDN | Indicates the Full Qualified Domain name (FQDN) to be used with IP address for the vSAN File Server instance. |

**Properties for vSAN Enabled Host**

The VMware Aria OperationsVMware Cloud Foundation Operations displays the following property for vSAN enabled host.

- Configuration|vSAN Enabled
- Configuration|vSAN|Encryption

**Properties for vSAN Cache Disk**

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the vSAN cache disk.

Properties for vSAN include:

| Component | Metrics |
|---|---|
| Configuration | • Configuration Properties\|Name<br>• Configuration Properties\|Size<br>• Configuration Properties\|Vendor<br>• Configuration Properties\|Type<br>• Configuration Properties\|Queue Depth<br>• Configuration\|vSAN\|Encryption<br>• Configuration \| Model |

*Table continued on next page*

*Continued from previous page*

| Component | Metrics |
|---|---|
| SCSI SMART Statistics | • SCSI SMART Statistics\|Media Wearout Indicator Threshold<br>• SCSI SMART Statistics\|Write Error Count Threshold<br>• SCSI SMART Statistics\|Read Error Count Threshold<br>• SCSI SMART Statistics\|Reallocated Sector Count Threshold<br>• SCSI SMART Statistics\|Raw Read Error Rate Threshold<br>• SCSI SMART Statistics\|Drive Temperature Threshold<br>• SCSI SMART Statistics\|Drive Rated Max Temperature Threshold<br>• SCSI SMART Statistics\|Write Sectors TOT Count Threshold<br>• SCSI SMART Statistics\|Read Sectors TOT Count Threshold<br>• SCSI SMART Statistics\|Initial Bad Block Count Threshold |

**Properties for vSAN Capacity Disk**

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the vSAN capacity disk.

Properties for vSAN include:

| Component | Metrics |
|---|---|
| Configuration | • Configuration Properties\|Name<br>• Configuration Properties\|Size<br>• Configuration Properties\|Vendor<br>• Configuration Properties\|Type<br>• Configuration Properties\|Queue Depth<br>• Configuration\|vSAN\|Encryption |
| SCSI SMART Statistics | • SCSI SMART Statistics\|Media Wearout Indicator Threshold<br>• SCSI SMART Statistics\|Write Error Count Threshold<br>• SCSI SMART Statistics\|Read Error Count Threshold<br>• SCSI SMART Statistics\|Reallocated Sector Count Threshold<br>• SCSI SMART Statistics\|Raw Read Error Rate Threshold<br>• SCSI SMART Statistics\|Drive Temperature Threshold<br>• SCSI SMART Statistics\|Drive Rated Max Temperature Threshold<br>• SCSI SMART Statistics\|Write Sectors TOT Count Threshold<br>• SCSI SMART Statistics\|Read Sectors TOT Count Threshold<br>• SCSI SMART Statistics\|Initial Bad Block Count Threshold |

**Properties for vSAN File Server**

The VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for vSAN file server.

- Configuration | vSAN | Primary
- Configuration | vSAN | FQDN

### Properties for vSAN File Share

The VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for vSAN file share.

- Configuration |vSAN| Domain Name
- Configuration | vSAN| Hard Quota
- Configuration |vSAN| Soft Quota
- Configuration |vSAN | Label|<key>
- Configuration |vSAN | Access Point|<key>
- Configuration | vSAN | Permission:<permission> | Client IP Range
- Configuration | vSAN | Permission:<permission> | Root Squash

### Properties for vSAN Storage Pool

The VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for vSAN Storage Pool.

| Property Name | Description |
|---|---|
| Configuration\|Number of Disks | Displays the total number of vSAN ESA disks in the storage pool. |

### Properties for vSAN ESA Disk

The VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the vSAN ESA Disk.

| Property Name | Description |
|---|---|
| Configuration\|Model | Displays the model number of the SCSI device. |
| Configuration\|Name | Displays the user configurable name for the SCSI device. |
| Configuration\|Queue Depth | Displays the queue depth of the SCSI device. |
| Configuration\|Size (GB) | Displays the size of SCSI device using the Logial Block Addressing Scheme (number of blocks) x (size of blocks). |
| Configuration\|Type | Displays the type of the SCSI device. |
| Configuration\|Vendor | Displays the vendor for the SCSI device. |
| Configuration \| vSAN \| Encryption | Indicates whether data encryption is enable on vSAN disk. If enabled all the VM data residing on the vSAN disk is encrypted. |

### Properties for Certificate Monitoring

VMware Aria OperationsVMware Cloud Foundation Operations displays the following certificate summary properties.

**Table 590: Adapter Instance Certificate Summary Properties, published on Adapter Instance Object**

| Property Name | Property Key | Description |
|---|---|---|
| End Date | Certificate Summary:endpointIdentifier\| End Date | End date of the adapter certificate. |

*Table continued on next page*

*Continued from previous page*

| Property Name | Property Key | Description |
|---|---|---|
| Start Date | Certificate Summary: endpointIdentifier \|Start Date | Start date of the adapter certificate. |
| Issuer DN | Certificate Summary:endpointIdentifier\| Issuer DN | Distinguished name of the adapter certificate issuer. |
| No. of days to expire | Certificate Summary:endpointIdentifier\| No. of days to expire | Number of days left before the expiration of the adapter certificate. |

**Table 591: Authentication Source Certificate Summary Properties, published on Universe Object**

| Property Name | Property Key | Description |
|---|---|---|
| End Date | Certificate Summary\|Authentication Sources:authenticationSourceId\|End Date | End date of the authentication source certificate. |
| Start Date | Certificate Summary\|Authentication Sources: authenticationSourceId\|Start Date | Start date of the authentication source certificate. |
| Issuer DN | Certificate Summary\|Authentication Sources:authenticationSourceId\|Issuer DN | Distinguished name of the authentication source certificate issuer. |
| No. of days to expire | Certificate Summary\|Authentication Sources:authenticationSourceId\|No. of days to expire | Number of days left before the expiration of the authentication source certificate |

**Table 592: Outbound Plugin Certificate Summary Properties, published on Universe Object**

| Property Name | Property Key | Description |
|---|---|---|
| End Date | Certificate Summary\|Outbound Plugins:outboundPluginId\|End Date | End date of the outbound plugin certificate |
| Start Date | Certificate Summary\|Outbound Plugins : outboundPluginId \|Start Date | Start date of the outbound plugin certificate |
| Issuer DN | Certificate Summary\| Outbound Plugins : outboundPluginId \|Issuer DN | Distinguished name of the outbound plugin certificate issuer. |
| No. of days to expire | Certificate Summary\| Outbound Plugins : outboundPluginId \|No. of days to expire | Number of days left before the expiration of the outbound plugin certificate |

### Properties for VMware Aria Automation

VMware Aria OperationsVMware Cloud Foundation Operations displays properties for VMware Aria Automation objects.

Some of the useful properties for project objects deployed through VMware Aria Automation are as follows:

- Project|CustomProperties: Custom properties defined for the project.
- Project|OrganizationID: Organization ID of the project.
- Project|userEmail: Email address of the user for the project.

One of the useful properties for the deployment object is:

- Deployment|User: User associated with the deployment.

One of the useful properties for the cloud zone object is:

- CloudAutomation|ResourceTags: Resource tags associated with the cloud zone.

One of the useful properties for the blueprint object is:

- Blueprint|User: User associated with the blueprint.

One of the useful properties for the CASworkd object is:

- CASWorld|metering|MeteringPolicyId: Metering policy ID associated with the CAS World object.

One of the useful properties for the virtual machine object is:

- Cloud Automation|CustomProperties: Custom properties associated with the virtual machine.

One of the useful properties for Cloud Zone is:
- Cloud Automation|Resource Tags: Resources tags associated with the cloud automation.

**Properties in the NSX**

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the NSX adapter.

**Table 593: Properties in the NSX Adapter**

| Resource | Properties common in NSX and NSX on VMware Cloud on AWS | Properties in NSX on-premise | Properties NSX on VMware Cloud on AWS |
|---|---|---|---|
| Management Cluster | | • NSXT Product Version<br>• Status Summary\|Cluster Status\| Management Cluster Status<br>• Status Summary\|Cluster Status\| Controller Cluster Status<br>• Status Summary\|vIDM Connection Status<br>• Status Summary\|Compute Managers\| <ComputeManagerName>\|Status<br>• Configuration Maximums<br>– Compute Manager count<br>– Prepared vC Cluster count | |
| Firewall Section | Summary<br>• Create Time<br>• Create User<br>• Last Modified Time<br>• Last Modified User<br>• Protection<br>• Revision<br>• System Owned<br>Configuration<br>• Firewall Rule Count Size | Configuration<br>• Firewall Stateful | Configuration<br>• Type<br>• Domain id<br>• Precedence<br>• Category |
| Transport Node | | • Summary | |

*Table continued on next page*

*Continued from previous page*

| Resource | Properties common in NSX and NSX on VMware Cloud on AWS | Properties in NSX on-premise | Properties NSX on VMware Cloud on AWS |
|---|---|---|---|
| **NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | – Create Time<br>– Create User<br>– Last Modified Time<br>– Last Modified User<br>– Protection<br>– Revision<br>– System Owned<br>– Summary\|FQDN<br>• Status Summary<br>– Transport Node State<br>– Transport Node Deployment State<br>– LCA Connectivity Status<br>– Management Plane Connectivity Status<br>– Host Node Deployment Status<br>– Management connection Status<br>– Controller connection Status<br>• Load Balancer Usage<br>– Current Small LB services<br>– Current Medium LB services<br>– Current Large LB services<br>– Current Extra Large LB services<br>– Current LB Pools<br>– Current LB Pool Members<br>– Current LB Virtual Servers<br>– Remaining Small LB services<br>– Remaining Medium LB services<br>– Remaining Large LB services<br>– Remaining Extra Large LB services<br>– Remaining LB Pool Members<br>• Tunnels\|<Tunnel-Name>\|Status<br>• File Systems\|<FileSystemMount><br>– Total<br>– Type<br>– File System ID | |
| Load Balancer Service | | • Summary<br>– Create Time<br>– Create User<br>– Last Modified Time<br>– Last Modified User<br>– Protection | |

*Table continued on next page*

*Continued from previous page*

| Resource | Properties common in NSX and NSX on VMware Cloud on AWS | Properties in NSX on-premise | Properties NSX on VMware Cloud on AWS |
|---|---|---|---|
| **NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | – Revision<br>– System Owned<br>– LB Service Operational Status | |
| Load Balancer Virtual Server<br>**NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | • Summary<br>– Create Time<br>– Create User<br>– Last Modified Time<br>– Last Modified User<br>– Protection<br>– Revision<br>– System Owned<br>– LB Virtual Operational State | |
| Load Balancer Pool | | • Summary<br>– Create Time<br>– Create User<br>– Last Modified Time<br>– Last Modified User<br>– Protection<br>– Revision<br>– System Owned<br>– Status | |

*Table continued on next page*

*Continued from previous page*

| Resource | Properties common in NSX and NSX on VMware Cloud on AWS | Properties in NSX on-premise | Properties NSX on VMware Cloud on AWS |
|---|---|---|---|
| **NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | | |
| Transport Zone<br>**NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | Summary<br>• Create Time<br>• Create User<br>• Last Modified Time<br>• Last Modified User<br>• Protection<br>• Revision<br>• Switch Mode<br>• System Owned | |
| Logical Router | • Summary<br>– Create Time<br>– Create User<br>– Last Modified Time<br>– Last Modified User<br>– Protection<br>– Revision<br>– System Owned | • Configuration<br>– Failover Mode<br>– High Availability Mode<br>– Edge Cluster Id<br>– Router Type<br>• Services Enabled<br>– HA Status Per Transport Node\|<TransportNodeID>\|HA Status<br>– Firewall Enabled<br>– Load balancer Enabled<br>– DNS Enabled<br>– L2VPN Enabled<br>– IPSEC VPN Enabled | |

*Table continued on next page*

*Continued from previous page*

| Resource | Properties common in NSX and NSX on VMware Cloud on AWS | Properties in NSX on-premise | Properties NSX on VMware Cloud on AWS |
|---|---|---|---|
| Router Service | 1. Tier-0 Router Services → BGP Service<br>– Summary\|BGP Neighbor Count<br>2. Tier-1 Router Services → NAT Rules<br>– Summary\|NAT Rule Count<br>3. Tier-1 Router Services → Static Routes<br>– Summary\|Static Route Count | • All logical routers → Static Routes → Summary\|Static Route Count<br>• All logical routers → NAT Rule → Summary\|NAT Rule Count<br>• Tier 0 → BGP Service → Summary<br>– ECMP Status<br>– Status<br>• Tier 0 → BFD Service → Summary<br>– Status<br>– BFD Neighbor Count<br>• Tier 0 → Route Redistribution → Summary<br>– Status<br>– Redistribution Rule count<br>• Tier 1 → Route Advertisement → Summary\|<br>– Route Advertisement Count<br>– Status | |
| Logical Switch | • Summary<br>– Create Time<br>– Create User<br>– Last Modified Time<br>– Last Modified User<br>– Protection<br>– Revision<br>– System Owned | • Summary<br>– Logical Switch State<br>• Configuration<br>– Replication Mode<br>– Admin State<br>– VNI | Configuration<br>• Type |
| Management Appliances<br>**NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | NSXT API Version | |
| Manager Node | | • NSXT Manager Node Version | |

*Table continued on next page*

*Continued from previous page*

| Resource | Properties common in NSX and NSX on VMware Cloud on AWS | Properties in NSX on-premise | Properties NSX on VMware Cloud on AWS |
|---|---|---|---|
| **NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | • Connectivity Status\|Management Plane Connectivity Status | |
| Group | Configuration Maximums\|Count<br>• IP Address Count<br>• Expressions Count<br>• vm Count | Configuration Maximums\|Count\|Tag Count | |
| Edge Cluster<br>**NOTE** This object is specific to NSX on-premise and is not available in NSX on VMware Cloud on AWS. | | Summary<br>• Create Time<br>• Create User<br>• Last Modified Time<br>• Last Modified User<br>• Protection<br>• Revision<br>• System Owned<br>• Edge Cluster Member Type | |

### Placement Group Properties

The following properties are available for each Placement Group instance in your VMware Aria OperationsVMware Cloud Foundation Operations environment.

**Table 594: Placement Group Properties**

| Service | Property |
|---|---|
| Placement Group | State |

*Table continued on next page*

*Continued from previous page*

| Service | Property |
|---|---|
| | Strategy |

### Properties for VeloCloud Gateway

VMware Aria OperationsVMware Cloud Foundation Operationsdisplays properties of VeloCloud Gateway objects.

Some of the useful properties for VeloCloud Gateway are as follows:
- Summary | Core Count
- Summary | Gateway Activation Status
- Summary | Gateway Network Interface Errors
- Summary | Gateway Time Zone
- Summary | ICMP Status
- Summary | Is Eth0 DPDK Enabled
- Summary | is Eth1 DPDK Enabled
- Summary | Registration Status
- Summary | VCO IP
- Summary | Version

### Properties for VeloCloud Orchestrator

VMware Aria OperationsVMware Cloud Foundation Operationsdisplays properties of VeloCloud Orchestrator objects.

Some of the useful properties for VeloCloud Orchestrator are as follows:
- General | DR SSH Tunnel Status
- General | Internet Connectivity
- General | IP Address
- General | NTP Time Zone

### Sustainability Properties

In the computed metrics, constant values for carbon dioxide emission and cost of power and tree offset for carbon dioxide emission have been used.

The values are taken from the following links:
- CO2 emission per kWh = 0.709 kg. Reference values from Greenhouse Gas Equivalencies Calculator
- Cost of Power = $0.108 per kWh. Based on the contiguous US average value. Reference from VMware TCO Reference Calculator .
- Tree offset for CO2 Emission = 16.511 kg (36.4 pounds of carbon per tree). Refer to Greenhouse Gases Equivalencies Calculator which is equivalent to 36.4/2.2046 kg of carbon per tree.
- Power consumption of a small server (1 socket, 10 cores, 32 GB RAM) = 0.1 kW (Assumption).

**Table 595:**

| Property Name | Object Type | Type | Default Value |
|---|---|---|---|
| Electricity Rate | Cluster Compute Resource | Numeric | 0.108 |
| CO2 Emission | Cluster Compute Resource | Numeric | 0.709 |
| Trees to Offer | Cluster Compute Resource | Numeric | 16.511 |

### Properties for Synthetic Monitoring

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the Business Application to show the state of Synthetic Monitoring.

**Properties for Synthetic Monitoring**

| Property Name | Property Key | Description |
|---|---|---|
| Synthetic Monitoring Activated | Synthetic Monitoring\| syntheticMonitoringActivated | Displays if Synthetic Monitoring is activated for the Business Application. |
| Synthetic Monitoring Configured | Synthetic Monitoring\| syntheticMonitoringConfigured | Displays if Synthetic Monitoring is configured for the Business Application. |

### Properties for Synthetic Monitoring Endpoints

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the Synthetic Monitoring Endpoints.

Properties for Synthetic Monitoring Endpoints

| Property Name | Property Key | Description |
|---|---|---|
| API Endpoint | Synthetic Monitoring:<api_id>\|API Endpoint | Displays if the API endpoint is enabled. |
| Request Type | Synthetic Monitoring:<api_id>\| Request Type | Displays the request type of the Synthetic Monitoring endpoints. |

### Properties for Policies

VMware Aria OperationsVMware Cloud Foundation Operations displays the following properties for the object type 'Policy'.

| Property Key | Property Name | Description |
|---|---|---|
| IsDefaultPolicy | Is Default Policy | This property indicates if a specific policy is the default policy. This property value is False for all policies except the default policy. |
| PolicyPriority | Policy Priority | This property indicates the priority of a specific policy. |

# VMware Aria Operations API Programming Guide (8.18)

The VMware Aria OperationsVMware Cloud Foundation Operations API Programming Guide provides information about the VMware Aria OperationsVMware Cloud Foundation Operations REST APIs, including how to use the REST API resources, authenticate, and construct REST API calls.

## Intended Audience

This information is intended for administrators and programmers who want to configure and manage VMware Aria OperationsVMware Cloud Foundation Operations programmatically using the VMware Aria OperationsVMware Cloud Foundation Operations REST API. The guide focuses on common use cases.

## Understanding the VMware Aria OperationsVMware Cloud Foundation Operations API

Developers can use the API to build interactive clients of VMware Aria OperationsVMware Cloud Foundation Operations. The API follows the REST style and is available to all licensed users.

VMware Aria OperationsVMware Cloud Foundation Operations clients communicate with the server over HTTP, exchanging representations of VMware Aria OperationsVMware Cloud Foundation Operations objects. These representations take the form of JSON or XML elements. You use HTTP GET requests to retrieve the current representation of an object, HTTP POST and PUT requests to create or modify an object, and HTTP DELETE requests to delete an object.

## How the VMware Aria OperationsVMware Cloud Foundation Operations API Works

You use a Web browser to communicate with the VMware Aria OperationsVMware Cloud Foundation Operations analytics engine, either through the product user interface or through API calls.

**Figure 21: VMware Aria OperationsVMware Cloud Foundation Operations Simplified Architecture**

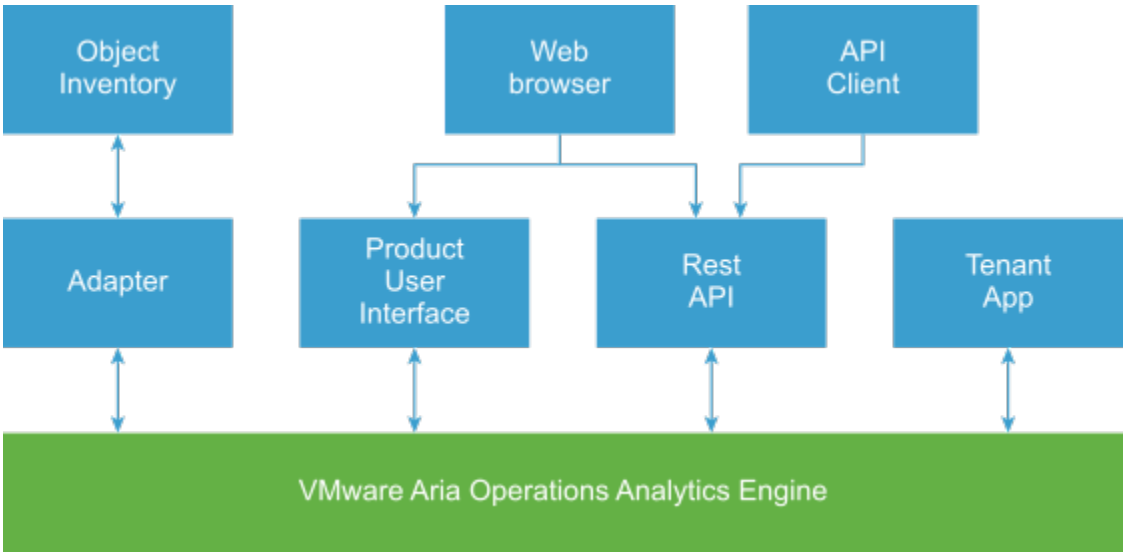


The adapter instance collects data from objects in your monitored environment. The VMware Aria OperationsVMware Cloud Foundation Operations analytics engine processes the data and displays the complete model in the graphical interface.

### Why Use the API

The API is most useful when there is a need to automate a well-defined workflow, such as repeating the same tasks to configure access control for new VMware Aria OperationsVMware Cloud Foundation Operations users. The API is also useful when performing queries on the VMware Aria OperationsVMware Cloud Foundation Operations data repository, such as retrieving data for particular assets in your virtual environment. In addition, you can use the API to extract all data from the VMware Aria OperationsVMware Cloud Foundation Operations data repository and load it into a separate analytics system.

### VMware Aria OperationsVMware Cloud Foundation Operations Terminology

The XML syntax you use to describe the objects for an adapter corresponds to the API code syntax but differs from what you find in the user interface. The following terms appear in the user interface. Included with the description of each term is the corresponding XML syntax used in an API call.

| | |
|---|---|
| Adapter types | Defines the adapter used to discover particular object types. For example, the vCenter adapter discovers objects connected to vSphere datacenters. The AWS adapter discovers AWS services and objects.<br>XML syntax: `adapterkinds`. |
| Object types | The class of entities that represent objects or information sources. Objects report data to the VMware Aria OperationsVMware Cloud Foundation Operations analytics engine. Virtual machines, datastores, and host systems are examples of object types defined in a vCenter adapter model.<br>XML syntax: `resourcekinds`. |

## Client Workflow Overview

VMware Aria OperationsVMware Cloud Foundation Operations API clients implement a REST workflow, making HTTP requests to the server and retrieving the information they need from the server’s responses.

### About REST

REST, an acronym for Representational State Transfer, describes an architectural style characteristic of programs that use the Hypertext Transfer Protocol (HTTP) to exchange serialized representations of objects between a client and a server. In the VMware Aria OperationsVMware Cloud Foundation Operations API, these representations are JSON or XML documents.

In a REST workflow, representations of objects are passed back and forth between a client and a server with the explicit assumption that neither party need know anything about an object other than what is presented in a single request or response. The URLs at which these documents are available often persist beyond the lifetime of the request or response that includes them.

### REST API Workflows

Application programs written to use a REST API use HTTP requests that are often executed by a script or other higher-level language to make remote procedure calls that create, retrieve, update, or delete objects that the API defines. In the VMware Aria OperationsVMware Cloud Foundation Operations REST API, these objects are defined by a collection of XML schemas. The operations themselves are HTTP requests, and so are generic to all HTTP clients.

To write a REST API client application, you must understand only the HTTP protocol, and the semantics of JSON or XML, the transfer format that the VMware Aria OperationsVMware Cloud Foundation Operations API uses. To use the API effectively in such a client, you must become familiar with the following concepts.

- The set of objects that the API supports, and what they represent.
- How the API represents these objects.

- How a client refers to an object on which it wants to operate.

The API reference includes a complete list of API requests. See Accessing Swagger Documentation for Schema Reference.

**VMware Aria OperationsVMware Cloud Foundation Operations API REST Requests**

To retrieve object representations, clients make HTTP requests to object references.

**Security**

The HTTP link between an API client and server is established using SSL. API clients configure token-based authentication to communicate with the server.

**Figure 22: Scenario: Provide user credentials to obtain details about alert with ID 123**

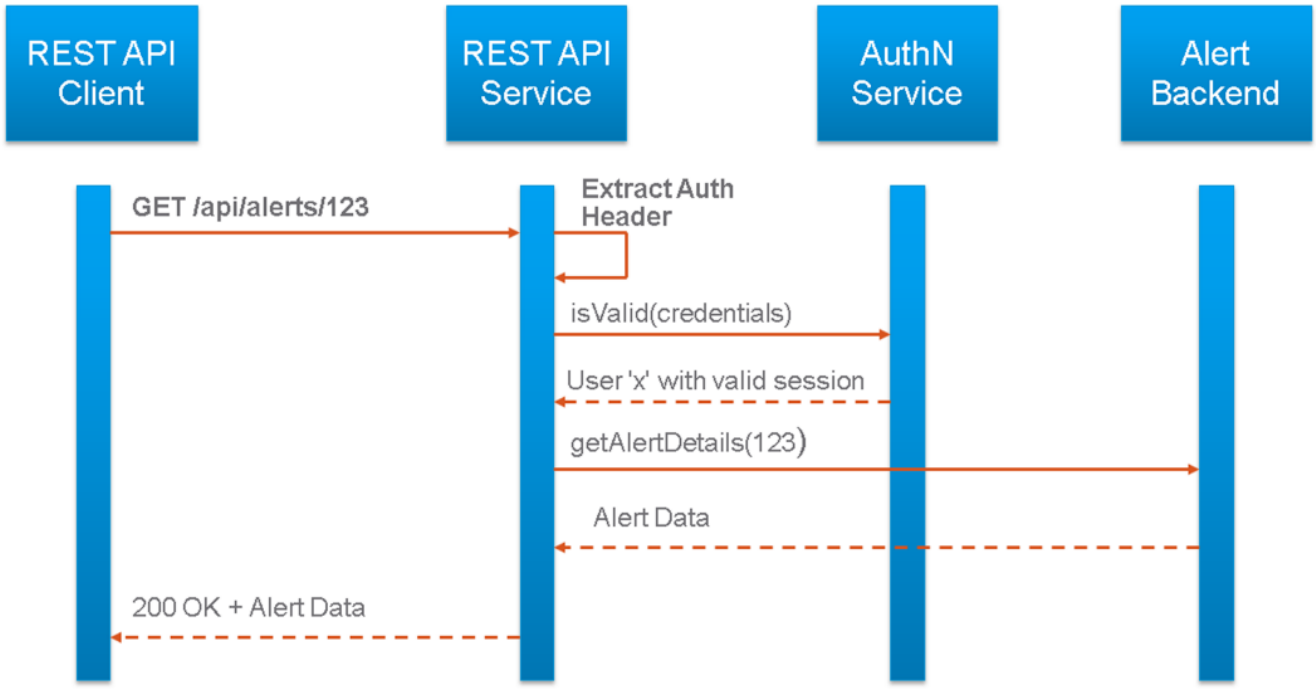


With token-based authentication, you POST a login request to the VMware Aria OperationsVMware Cloud Foundation Operations API server, supplying valid user credentials to obtain an authentication token. The following example presents a token-based authentication scenario.

1. You obtain valid user credentials for your VMware Aria OperationsVMware Cloud Foundation Operations instance. This is applicable for non-SSO authenticaton.
2. POST a request to the REST endpoint for authentication.
   `https://RESTendpoint.example.com/suite-api/api/auth/token/acquire`

   The request body includes the user name, password, and authentication source.
3. In the response body, the endpoint returns the token, expiry date, and time.
4. For further communication, you include the token object in the Authorization header with the format :
   `Authorization: OpsToken <vROps_token>`

**NOTE**
The old format,
`Authorization: vRealizeOpsToken <vROps_token>`
continues to be supported in VMware Aria Operations.

If you acquired the token externally from an SSO source (without using `/suite-api/api/auth/token/acquire API`), the Authorization header is of the format:
`Authorization: SSO2Token <SSO_SAML_TOKEN>`

5. You can invalidate the token before the expiration date and time by sending a POST request to the logout endpoint.
   `POST https://RESTendpoint.example.com/suite-api/api/auth/token/release`

You can call the VMware Cloud Foundation Operations APIs after generating the Cloud Services authentication token. Use the following endpoint:
`https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api`

**NOTE**
If your organization is located outside of the United States, use the country abbreviation for your API endpoint. Using the wrong endpoint will result in a 404 error. For example, if your organization is located in Australia, your country abbreviation is AU and your API endpoint is:
`https://au.www.mgmt.cloud.vmware.com/vrops-cloud/suite-api`

You must pass the Cloud Services authentication token with every request as an HTTP header in the following format:
`Authorization: CSPToken {CSP Auth Token}`

For example:
```
curl -k https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/resources -H "Content-Type: application/json" -H "Accept: application/json" -H "Authorization: CSPToken abc12345...
```

**Request Headers**

The following HTTP headers are typically included in API requests:

| | |
|---|---|
| Accept-Language | To specify the language desired in responses, use the `Accept-Language` request header. Message strings in `ErrorType` responses are localized. To request a response with message strings localized to French, use the following header:<br>`Accept-Language: fr-FR` |
| Authorization | All requests to create an API session must include an `Authorization` header of the form prescribed by the identity provider that your organization uses |
| Content-Type | Requests that include a body must include an appropriate HTTP `Content-Type` header.<br>• For a request body in XML, the header must include `Content-Type: application/xml`<br>• For a request body in JSON, the header must include `Content-Type: application/json` |
| Accept | To specify the desired response format , include the `Accept` request header. |

*Table continued on next page*

*Continued from previous page*

| | |
|---|---|
| | • For a response in XML, the header must include `Accept: application/xml`<br>• For a response in JSON, the header must include `Accept: application/json` |

**Request Bodies in XML**

For a request body written in XML, VMware Aria OperationsVMware Cloud Foundation Operations uses a validating XML parser that requires elements in a request body to agree with the schema in order and number. Request bodies are rejected as invalid unless they meet the following criteria:

- XML namespace attributes must be supplied for all namespaces represented by elements in the request.
- If multiple namespaces are represented in the request, XML namespace attributes must include an identifying prefix, and that prefix must be used with all elements from that namespace.
- All required elements must appear in request bodies. All elements that appear in request bodies must appear in the order that the schema establishes, and with content that conforms to the type constraint that the schema specifies.

### VMware Aria OperationsVMware Cloud Foundation Operations API REST Responses

All responses include an HTTP status code and, unless the status code is 204 (No Content), an Accept header. Response content depends on the request. Some responses include a document body, some include only a URL, and some are empty.

**HTTP Response Codes**

An API client can expect a subset of HTTP status codes in a response.

**Table 596: HTTP Status Codes that the API Returns**

| Status Code | Status Description |
|---|---|
| `200 OK` | The request is valid and was completed. The response includes a document body. |
| `201 Created` | The request is valid. The requested object was created and can be found at the URL specified in the Location header. |
| `202 Accepted` | The request is valid and a task was created to handle it. This response is usually accompanied by a `TaskStatus` element . |
| `204 No Content` | The request is valid and was completed. The response does not include a body. |
| `400 Bad Request` | The request body is malformed, incomplete, or otherwise invalid. |
| `401 Unauthorized` | Login failed or authentication token has expired. |
| `403 Forbidden` | The user is not authenticated or does not have adequate privileges to access one or more objects specified in the request. |
| `404 Not Found` | The object specified in the request could not be found. |
| `405 Method Not Allowed` | The HTTP method specified in the request is not supported for this object. |

*Table continued on next page*

*Continued from previous page*

| Status Code | Status Description |
|---|---|
| `406 Not Acceptable` | The resource identified by the request is not capable of generating a response of the type specified in the request's `Accept` header. |
| `415 Unsupported Media Type` | The resource identified by the request does not support a request of the specified `Content-Type` and HTTP method. |
| `422 Not Found` | Usually indicates a malformed request URL or request body. |
| `429 Too Many Requests` | A client has sent too many requests or multiple clients are sending too many simultaneous requests and the server is unable to process them due to rate limits.<br>To work around the problem, try sending the request again later. |
| `500 Internal Server Error` | The request was received but could not be completed because of an internal error on the server. |
| `503 Service Unavailable` | The server is currently unable to handle the request due to a temporary condition such as resource exhaustion or server maintenance. |
| `504 Gateway Timeout` | The server, while acting as a gateway or proxy, did not receive a timely response from the upstream server specified by the request URL. |

## Using the API with VMware Aria OperationsVMware Cloud Foundation Operations

You can use a browser or an HTTP client program to send requests and receive responses.

### REST Client Programs

Any client application that can send HTTPS requests is an appropriate tool for developing REST applications with the VMware Aria OperationsVMware Cloud Foundation Operations API. REST client plug-ins are available for most browsers and many IDEs. The following open-source programs are commonly used:

- cURL. http://curl.haxx.se
- Postman application. http://www.getpostman.com

In addition, VMware provides language-specific client bindings for the VMware Aria OperationsVMware Cloud Foundation Operations API. See Accessing Swagger Documentation for Schema Reference.

### Accessing Swagger Documentation for Schema Reference

The VMware Aria OperationsVMware Cloud Foundation Operations REST API documentation includes reference material for all elements, types, queries, and operations in the VMware Aria OperationsVMware Cloud Foundation Operations API. It also includes the schema definition files.

Swagger based API documentation is available with the product, with the capability of making REST API calls right from the landing page.

To access the API documentation, you must first log into VMware Cloud Foundation Operations at the URL of your VMware Cloud Foundation Operations instance.

For example, if the URL of your instance is `https://operations.example.comhttps://www.mgmt.cloud.vmware.com`, the API reference is available from: `https://operations.example.com/suite-api/doc/swagger-ui.html`https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/doc/swagger-ui.html.

> **NOTE**
> If your organization is located outside of the United States, use the country abbreviation for your API endpoint. Using the wrong endpoint will result in a 404 error. For example, if your organization is located in Australia, your country abbreviation is AU and your API documentation will be available at:
> `https://au.www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/doc/swagger-ui.html`

Language-specific client bindings are available from:

`https://operations.example.com/suite-api/`

`https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/`

**About the VMware Aria OperationsVMware Cloud Foundation Operations API Examples**

All examples include HTTP requests and responses. These examples show the workflow and content associated with operations such as creating and querying for information about objects in your monitored environment.

Example request bodies are in JSON. Request headers required by the VMware Aria OperationsVMware Cloud Foundation Operations API are included in example requests that are not fragments of a larger example.

Most example responses show only those elements and attributes that are relevant to the operation being discussed. Ellipses (...) indicate omitted content within response bodies.

## Getting Started with the API

API clients and VMware Aria OperationsVMware Cloud Foundation Operations servers communicate over HTTPS, exchanging XML representations of API objects.

This simple example of a REST API workflow shows how to obtain a list of all metrics for a virtual machine object type that is included in the model definition of the VMware vCenter® adapter. Using the API, you can obtain the complete list of available metrics for any object type.

## Acquire an Authentication Token

VMware Cloud Foundation Operations requires API requests to be authenticated. The first step in this workflow is to obtain an authentication token.

- Secure a channel between the web browser and the VMware Cloud Foundation Operations server. Open a browser and enter the URL of a VMware Cloud Foundation Operations instance such as:
  `https://vrealize.example.com`

  The system warns that your connection is not private. Click through to confirm the security exception and establish an SSL handshake.
- Verify that you can access the APIs. Enter the URL of your VMware Cloud Foundation Operations instance with `suite-api/docs/rest/index.html` added to the end, such as:
  `https://vrealize.example.com/suite-api/docs/rest/index.html`
- Verify that you have the login credentials for a user of your VMware Cloud Foundation Operations instance.

To obtain an authentication token, the login request supplies the user credentials in a form that Basic HTTP authentication requires. In this example, the user is logging in to a VMware Cloud Foundation Operations instance with URL `https://vrealize.example.com/`.

**NOTE**
This example uses token-based authentication. For more information regarding authentication mechanisms, see Security.

Using `authSource`, you can import and authenticate users and user group information that reside on another machine. For example, you can authenticate users from LDAP, Active Directory, VMware Identity Manager, Single Sign-On and so on. When you import user account information that resides on another machine, you must define the criteria used to import the user accounts from the source machine.
After creating an auth source you can use it for acquiring a token by specifying the name. The default auth source type is LOCAL.

1. POST a request to the login URL to acquire a token.

   See Login Request and Response.

   1. You obtain valid user credentials for your VMware Aria OperationsVMware Cloud Foundation Operations instance.
   2. POST a request to the REST endpoint for authentication (non-SSO).
      `https://RESTendpoint.example.com/suite-api/api/auth/token/acquire`

      The request body includes the user name, password, and authentication source.
   3. In the response body, the endpoint returns the token, expiry date, and time.
   4. For further communication, you include the token object in the Authorization header with the format :
      `Authorization: OpsToken <vROps_token>`

      **NOTE**
      The old format,
      `Authorization: vRealizeOpsToken <vROps_token>`

      continues to be supported in VMware Aria Operations.

      If you acquired the token externally from an SSO source (without using `/suite-api/api/auth/token/acquire API`), the Authorization header is of the format:
      `Authorization: SSO2Token <SSO_SAML_TOKEN>`

2. Examine the response.

   A successful request returns an `ops` authorization token, which you must included in subsequent API requests.

## Login Request and Response

This example shows a request and response for a user with the login username: `vRealize-user` and password: `vRealize-dummy-password`.

Request header:

```
POST https://vrealize.example.com/suite-api/api/auth/token/acquire

Content-Type: application/json

Accept: application/json
```

Request body in JSON format:

```
{

  "username" : "vRealize-user",
```

```
  "password" : "vRealize-dummy-password"
}
```

Response in JSON:

```
200 OK
{
  "token": "8f868cca-27cc-43d6-a838-c5467e73ec45::77cea9b2-1e87-490e-b626-e878beeaa23b",
  "validity": 1470421325035,
  "expiresAt": "Friday, August 5, 2016 6:22:05 PM UTC",
  "roles": []
}
```

The response code indicates whether the request succeeded, or how it failed.

- If the request is successful, the server return HTTP response code 200 (OK) and re-usable `ops` authorization token that expires after six hours. This token must be included in each subsequent API request.
- If the authorization header is missing for the request, the server returns HTTP response code 403.
- If the credentials supplied in the Authorization header are invalid, the server returns HTTP response code 401.

## Generate Cloud Services Authentication Tokens

Developers can use the API to build interactive clients of VMware Cloud Foundation Operations. The API follows the REST style and is available to all licensed users.

1. Generate the API tokens from the Cloud Services toolbar to authenticate yourself when you make authorized API connections. See, Generate API Tokens.
2. Call the Cloud Services API using the API token to generate the Cloud Services authentication token.

   ```
   curl -k -X POST "https://console.cloud.vmware.com/csp/gateway/am/api/auth/api-tokens/
   authorize" -H "accept: application/json" -H "Content-Type: application/x-www-form-
   urlencoded" -d "refresh_token={CSP API Token}”
   ```

   This returns a JSON string with the following structure in which `access_token` is the Cloud Services authentication token:

   ```
   {
      "id_token",
      "token_type",
      "expires_in",
      "scope",
      "access_token",
      "refresh_token”
   }
   ```

**NOTE**
The Cloud Services authentication token is valid only for 30 minutes after it is generated.

## Find the Adapter Type and Object Type

Your VMware Aria OperationsVMware Cloud Foundation Operations instance includes multiple adapter types. To find the adapter type for the vCenter adapter, you make a GET request to retrieve a list of all adapter types. The API response includes all the object types that the adapter monitors.

Verify that you are logged in to the VMware Aria OperationsVMware Cloud Foundation Operations instance.

1. Make a GET request for all adapter types.

   ```
   GET https://operations.example.com/suite-api/api/adapterkinds
   ```

   ```
   GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapterkinds
   ```

2. Examine the response to find the vCenter adapter and list of monitored object types.

   See the response portion of Determine the Adapter Type and Object Types for the vCenter Adapter.

## Determine the Adapter Type and Object Types for the vCenter Adapter

This example finds the adapter type for the vCenter adapter and all the object types included in the adapter model definition.

Request header:

```
GET https://operations.example.com/suite-api/api/adapterkinds
```

```
Content-Type: application/json
```

```
Authorization: OpsToken <vROps_token>
```

```
Accept: application/json
```

**NOTE**
The old format,
`Authorization: vRealizeOpsToken <vROps_token>`
continues to be supported in VMware Aria Operations.

```
GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapterkinds
```

```
Content-Type: application/json
```

```
Authorization: CSPToken <csp auth token>
```

```
Accept: application/json
```

Where *vROps_token<csp auth token>* is the token that you obtained from the response in Acquire an Authentication TokenGenerate Cloud Services Authentication Tokens.

Snippet of the response in JSON for the vCenter Adapter:

```
200 OK
{
  "key": "VMWARE",
  "name": "vCenter Adapter",
  "description": "Provides the connection information and credentials required...",
  "adapterKindType": "GENERAL",
  "describeVersion": 573,
  "identifiers": [],
  "resourceKinds": [
    "ClusterComputeResource",
    "ComputeResource",
    "CustomDatacenter",
    "Datacenter",
    "Datastore",
    "StoragePod",
    "DatastoreFolder",
    "VM Entity Status",
    "Folder",
    "HostFolder",
    "HostSystem",
    "NetworkFolder",
    "ResourcePool",
    "VMwareAdapter Instance",
    "VirtualMachine",
    "VMFolder",
    "DistributedVirtualPortgroup",
    "VmwareDistributedVirtualSwitch",
    "vSphere World"
  ],
  ...
}
```

For the vCenter adapter, the `adapter-kind` key is `VMWARE`. The `resourceKinds` are the object types that the vCenter adapter monitors. For virtual machine object type, the `resourceKinds` is `VirtualMachine`.

## Generate a List of All Metrics for the Object

To generate a complete list of metrics for any virtual machine defined in the vCenter adapter model, you make a GET request to the URL with the adapter type and the object type.

Verify that the following requirements are met:
- You are logged in to the VMware Aria OperationsVMware Cloud Foundation Operations instance.
- You know the `adapterKind` value for the vCenter adapter and the `resourceKinds` value for the virtual machine. See Determine the Adapter Type and Object Types for the vCenter Adapter

1. Make a GET request to obtain the metadata for metrics.

```
GET https://operations.example.com/suite-api/api/adapterkinds/VMWARE/resourcekinds/
VirtualMachine/statkeys
```

```
GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapterkinds/VMWARE/
resourcekinds/VirtualMachine/statkeys
```

2. Compare the metrics listed in the response to metrics displayed in the user interface. See Virtual Machine Metrics from the API and in the User Interface

## Virtual Machine Metrics from the API and in the User Interface

This example shows how the virtual machine metrics listed in the XML response compare to the metrics displayed in the VMware Aria OperationsVMware Cloud Foundation Operations user interface.

Request:

```
GET https://operations.example.com/suite-api/api/adapterkinds/VMWARE/resourcekinds/
VirtualMachine/statkeys

Content-Type: application/json

Authorization: OpsToken <vROps_token>

Accept: application/json
```

> **NOTE**
> The old format,
> `Authorization: vRealizeOpsToken <vROps_token>`
> continues to be supported in VMware Aria Operations.

```
GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapterkinds/VMWARE/
resourcekinds/VirtualMachine/statkeys

Content-Type: application/json

Authorization: CSPToken <csp auth token>

Accept: application/json
```

Where:
- `VMWARE` is the `adapterKindKey`.
- `VirtualMachine` is the `resourceKindKey`.

- *vROps_tokencsp auth token* is the token that you obtained from the response in Acquire an Authentication Token.Generate Cloud Services Authentication Tokens.

Snippet of the response in JSON:

```
200 OK


{
    "resourceTypeAttributes": [
        ...
        {
            "key": "mem|host_workload",
            "name": "Memory|Host Workload",
            "description": "Host Workload (%)",
            "defaultMonitored": false,
            "rollupType": "AVG",
            "instanceType": "INSTANCED",
            "unit": "%",
            "dataType2": "FLOAT",
            "monitoring": false,
            "property": false
        },
        ...
    ]
}
```

Every `resourceTypeAttribute` in the response is a metric with metadata for a virtual machine object. The `name` corresponds to text displayed in the VMware Aria OperationsVMware Cloud Foundation Operations user interface. In this example, the snippet lists metrics for Memory and Host Workload.

To compare metrics in the response with metrics in the user interface, log in to the VMware Aria OperationsVMware Cloud Foundation Operations instance running on `operations.example.commgmt.cloud.vmware.com` and navigate to the metrics for a virtual machine. The following example shows where you find metrics for Memory(Host) and Workload.

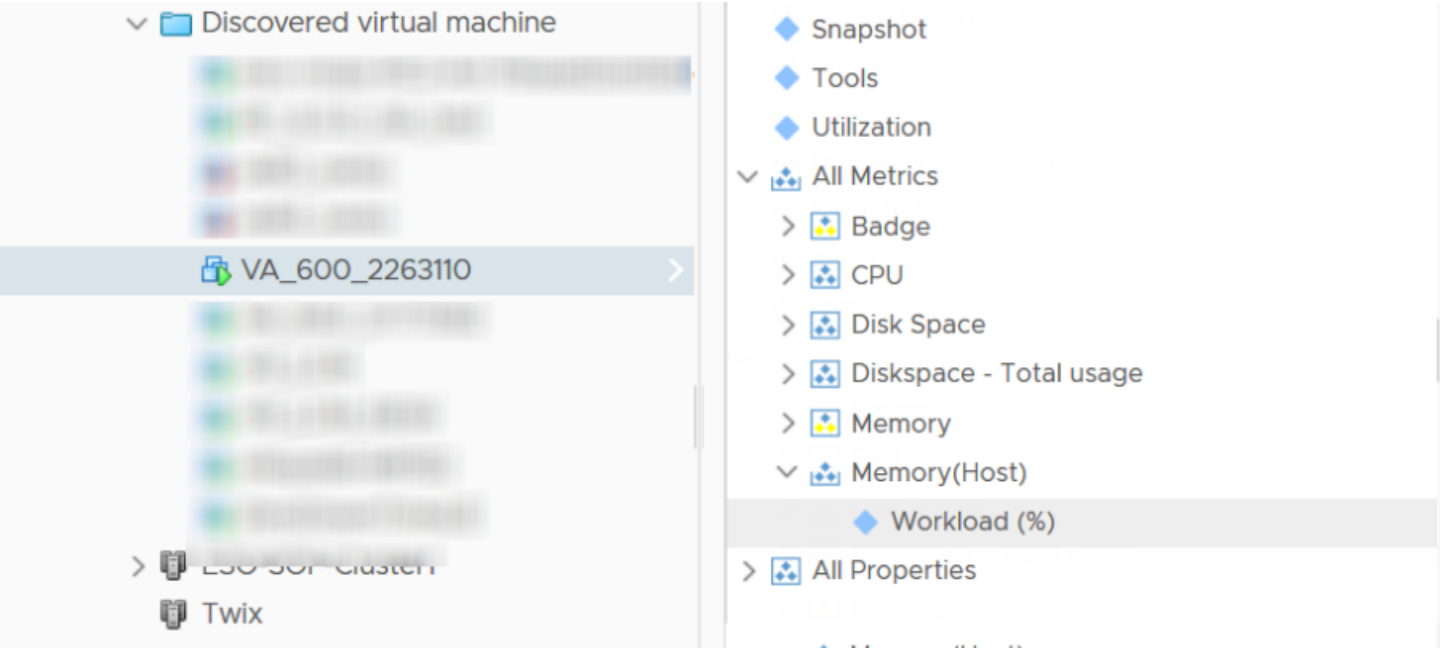


The example shows how to retrieve metrics for the virtual machine object type. To retrieve metrics for other object types, replace `VirtualMachine` in the GET request with other `resourceKinds`.

## Configuring an Adapter Instance

After installing a solution that includes a management pack with an adapter, you must configure an adapter instance to collect data from the objects in the adapter model definition. You can use the VMware Aria OperationsVMware Cloud Foundation Operations API to configure an adapter instance.

This use case example shows how to configure an adapter instance for a vSphere Solution and includes:

- summary of operations with request, request body, and response for each
- specific procedure for each operation

## Summary of Configuring an Adapter Instance Requests

You make sequential API requests to configure an adapter instance. Responses from earlier requests yield information required for a subsequent requests.

**Table 597: Summary of Requests**

| Operation | Request | Request Body | Response |
|---|---|---|---|
| Get all solutions registered with the product and identify the adapter types | GET <*API-URL*>/suite-api/api/solutions | None | `adapterkindkeys` |
| Get all the object types for a particular adapter type. | GET <*API-URL*>/suite-api/api/adapterkinds/{key}/resourcekinds | None | `resourceIdentifierTypes` |
| Create an adapter instance object | POST <*API-URL*>/suite-api/api/adapters | Values for `resourceIdentifiers` and `credential` | `adapterid` |

*Table continued on next page*

*Continued from previous page*

| Operation | Request | Request Body | Response |
|---|---|---|---|
| Patch an adapter instance to acknowledge the presented certificate | PATCH <*API-URL*>/suite-api/api/adapters | API response of POST <*API-URL*>/suite-api/api/adapters | API response of POST <*API-URL*>/suite-api/api/adapters without `adapter-cerfiticates` |
| Start adapter monitoring | PUT <*API-URL*>/suite-api/api/adapters/{adapterid}/monitoringstate/start | None | 200 OK |

## Identify the Solution and Its Adapters

Your VMware Aria OperationsVMware Cloud Foundation Operations instance may have several solutions installed. To find the vSphere solution and its adapter types, you make a GET request to retrieve a list of all solutions. The response includes all the adapters included with the solution.

Verify that you can log in to the API URL for a VMware Aria OperationsVMware Cloud Foundation Operations instance. See Acquire an Authentication Token.

For this example use case, the *API-URL* for the VMware Aria OperationsVMware Cloud Foundation Operations instance is `operations.example.commgmt.cloud.vmware.com`.

1. Make a GET request to list all the solutions.

   ```
   GET https://operations.example.com/suite-api/api/solutions
   ```

   ```
   GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/solutions
   ```

2. Examine the response to find the vSphere solution and its adapter types.

   See the response portion of Adapter Types for the vSphere Solution.

## Adapter Types for the vSphere Solution

This example lists all the installed solutions and the adapter types for each.

Request header:

```
GET https://operations.example.com/suite-api/api/solutions
```

```
GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/solutions
```

The response in JSON:

```
200 OK
```

```
{
```

```
  "solution":
```

```
    {
```

```
      "id": "MPforLogInsight",
```

```
      "name": "VMware vRealize Operations Management Pack for Log Insight",
```

```
      "version": "6.0.3171089",
```

```
      "description": "VMware vRealize Operations Management Pack for Log Insight... ",
```

```
      "vendor": "VMware Inc.",
      "adapterKindKeys": [
        "LogInsightAdapter"
      ]
    },
    {
      "id": "ep-ops-os-and-availability",
      "name": "Operating Systems / Remote Service Monitoring",
      "version": "1.0.4071095",
      "description": "The End Point Operations Management Solution for Operating... ",
      "vendor": "VMware Inc.",
      "adapterKindKeys": [
        "ep-ops-os-and-availability-kind"
      ]
    },
    {
      "id": "VMware vSphere",
      "name": "VMware vSphere",
      "version": "6.0.7496664",
      "description": "Manages vSphere objects such as Clusters, Hosts...",
      "vendor": "VMware Inc.",
      "adapterKindKeys": [
        "VMWARE",
        "PythonRemediationVcenterAdapter"
      ]
    }
  ]
}
```


The response shows three solutions installed:
- Management Pack for Log Insight solution
- End Point Operations solution
- vSphere solution

The vSphere solution has two adapter types:
- VMWARE

- PythonRemediationVcenterAdapter

For the vCenter adapter, the adapter type is VMWARE.

## Identify the Object Types Required for the Adapter

After you determine that you want to create an instance of the vCenter adapter, you must identify the required object types for that adapter. You make a GET request to retrieve a list of all object types for the vCenter adapter.

Verify that you know the adapter type for the vCenter adapter.

1. Make a GET request to list all the object types for the vCenter adapter.

   ```
   GET https://operations.example.com/suite-api/api/adapterkinds/VMWARE/resourcekinds
   ```

   ```
   GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapterkinds/VMWARE/
   resourcekinds
   ```

2. Examine the response to identify the required object types..

   See the response portion of Object Types Required for the vCenter Adapter.

## Object Types Required for the vCenter Adapter

This example finds all the object types for the vCenter adapter.

Request header:

```
GET https://operations.example.com/suite-api/api/adapterkinds/VMWARE/resourcekinds
```

```
GET https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapterkinds/VMWARE/
resourcekinds
```

Snippet of the response in JSON:

```
200 OK
```


```
{
  "key": "VMwareAdapter Instance",
  "name": "vCenter Server",
  "adapterKind": "VMWARE",
  "resourceKindType": "ADAPTER_INSTANCE",
  "resourceKindSubType": "NONE",
  "resourceIdentifierTypes": [
    {
      "name": "AUTODISCOVERY",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    {
      "name": "DISABLE_COMPUTATION_BASED_ON_CONSUMERS",
```

```
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    {
      "name": "DV_PORT_GROUP_DISABLED",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    {
      "name": "DVS_DISABLED",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    {
      "name": "PROCESSCHANGEEVENTS",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    {
      "name": "VCURL",
      "dataType": "STRING",
      "isPartOfUniqueness": true
    },
    ...
    {
      "name": "VM_LIMIT",
      "dataType": "INTEGER",
      "isPartOfUniqueness": false
    }
    ],
  ...
}
```

This snippet shows the Resource Kind with the attribute `"resourceKindType": "ADAPTER_INSTANCE"`. Any object type that has the resource identifier `"isPartOfUniqueness":true` requires a value for that object type with the API request to create the adapter instance.

An adapter instance of the vCenter adapter requires a value for `VCURL` or the URL of the vCenter.

## Create the Adapter Instance

After you identify the object types required for the adapter, you provide parameter values for the object types to create an adapter instance. Your POST request includes a request body with the required parameters.

Verify that you have an IP address and credentials for a vCenter.

To create an adapter instance, the VCURL setting is mandatory.

1. Make a POST request to create the adapter instance.

   ```
   POST https://operations.example.com/suite-api/api/adapters
   ```

   ```
   POST https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapters
   ```

2. Examine the response to find the name for the vSphere Solution and its adapter types.

   See the response portion of Adapter Instance.

## Adapter Instance

This example creates the adapter instance for a vCenter with the following parameters:

- Display Name: VC Adapter Instance
- Description: A vCenter Adapter Instance for VC 12.345.678.9
- vCenter Server IP address: https://12.345.678.9
- Credential name: VC-Credential-1
- User Name: administrator@vsphere.local
- Password: VC-dummy-passwd

`AUTODISCOVERY` and `PROCESSCHANGEEVENTS` are optional, but are included to show additional examples of resource identifiers in the request body and in the response.

Request header:

```
POST https://operations.example.com/suite-api/api/adapters
```

```
POST https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapters
```

Request body in JSON format:

```
{
  "name" : "VC Adapter Instance",
  "description" : "A vCenter Adapter Instance for VC 12.345.678.9",
  "collectorId" : "1",
  "adapterKindKey" : "VMWARE",
  "resourceIdentifiers" : [
    {
      "name" : "AUTODISCOVERY",
```

```
      "value" : "true"
    },
    {
      "name" : "PROCESSCHANGEEVENTS",
      "value" : "true"
    },
    {
      "name" : "VCURL",
      "value" : "https://12.345.678.9"
    }
  ],
  "credential" : {
    "id" : null,
    "name" : "VC-Credential-1",
    "adapterKindKey" : "VMWARE",
    "credentialKindKey" : "PRINCIPALCREDENTIAL",
    "fields" : [
      {
        "name" : "USER",
        "value" : "administrator@vsphere.local"
      },
      {
        "name" : "PASSWORD",
        "value" : "VC-dummy-passwd"
      }
    ],
  },
}
```

Snippet of the response in JSON:

```
201 Created
{
  "resourceKey": {
```

```
"name": "VC Adapter Instance",
"adapterKindKey": "VMWARE",
"resourceKindKey": "VMwareAdapter Instance",
"resourceIdentifiers": [
  {
    "identifierType": {
      "name": "AUTODISCOVERY",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": "true"
  },
  {
    "identifierType": {
      "name": "DISABLE_COMPUTATION_BASED_ON_CONSUMERS",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": ""
  },
  {
    "identifierType": {
      "name": "DV_PORT_GROUP_DISABLED",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": ""
  },
  {
    "identifierType": {
    "name": "DVS_DISABLED",
    "dataType": "STRING",
```

```
    "isPartOfUniqueness": false
    },
    "value": ""
  },
  {
    "identifierType": {
      "name": "PROCESSCHANGEEVENTS",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": "true"
  },
  {
    "identifierType": {
      "name": "VCURL",
      "dataType": "STRING",
      "isPartOfUniqueness": true
    },
    "value": "https://12.345.678.9"
  },
  {
    "identifierType": {
      "name": "VM_FOLDER_DISABLED",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": ""
  },
  {
    "identifierType": {
      "name": "VM_LIMIT",
      "dataType": "STRING",
```

```
          "isPartOfUniqueness": false
        },
        "value": ""
      }
    ]
  },
  "description": "A vCenter Adapter Instance for VC 12.345.678.9",
  "collectorId": 1,
  "collectorGroupId": "909c2fbf-2c2c-4957-9a75-21bf2a887d31",
  "credentialInstanceId": "65081a8d-d462-43b2-b4e0-596eaf3d497e",
  "monitoringInterval": 5,
  "adapter-certificates": [
    {
      "thumbprint": "2520fb4351bc91ee7b82ef7cc54a8d88fa893da9",
      "certificateDetails": "[
        Version: V3 Subject: C=US, CN=12.345.678.9
        Signature Algorithm: SHA256withRSA, OID = 1.2.840.113549.1.1.11
        Key: Sun RSA public key, 2048 bits modulus: ...
        Validity: [From: Wed Jul 15 19:26:51 UTC 2015, To: Tue Jul 08 11:26:30 UTC 2025]
        Issuer: O=W12R2UINanduVC, C=US, DC=local, DC=vsphere, CN=CA ...
          ...
      ]"
    }
  ],
  ...
  "id": "a97bd204-e3e5-404b-a219-e2b20cf158d2"
}
```


The API creates a new adapter with an internally generated UUID that uniquely identifies the object. The API response includes the certificates that vCenter 12.345.678.9 presents. The value of the adapter instance ID is used to start monitoring and collecting data.

## Provide Proof of Certificate Validity

Before VMware Aria OperationsVMware Cloud Foundation Operations can connect to the vCenter and start collecting data, it needs to verify that data sources discovered by the adapter instance are presenting valid certificates. Your PATCH

request provides the proof that the certificates are valid by including a request body that is the response from the POST request used to create the adapter.

Verify that you have the response from the POST request used to create the adapter. See the response in Adapter Instance.

1. Make a PATCH request to notify the system that the user has accepted the certificate presented by the vCenter.

   ```
   PATCH https://operations.example.com/suite-api/api/adapters
   ```

   ```
   PATCH https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapters
   ```

## Certificate Validation

In this example, the request body for the PATCH request is the same as the response from the POST request used to create the adapter instance.

Request header:

```
PATCH https://operations.example.com/suite-api/api/adapters
```

```
PATCH https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapters
```

Request body in JSON format:

```
{
  "resourceKey": {
    "name": "VC Adapter Instance",
    "adapterKindKey": "VMWARE",
    "resourceKindKey": "VMwareAdapter Instance",
    "resourceIdentifiers": [
      {
        "identifierType": {
          "name": "AUTODISCOVERY",
          "dataType": "STRING",
          "isPartOfUniqueness": false
        },
        "value": "true"
      },
      {
        "identifierType": {
          "name": "DISABLE_COMPUTATION_BASED_ON_CONSUMERS",
          "dataType": "STRING",
          "isPartOfUniqueness": false
```

```
    },
    "value": ""
  },
  {
    "identifierType": {
      "name": "DV_PORT_GROUP_DISABLED",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": ""
  },
  {
    "identifierType": {
    "name": "DVS_DISABLED",
    "dataType": "STRING",
    "isPartOfUniqueness": false
    },
    "value": ""
  },
  {
    "identifierType": {
      "name": "PROCESSCHANGEEVENTS",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": "true"
  },
  {
    "identifierType": {
      "name": "VCURL",
      "dataType": "STRING",
      "isPartOfUniqueness": true
```

```
      },
      "value": "https://12.345.678.9"
    },
    {
      "identifierType": {
        "name": "VM_FOLDER_DISABLED",
        "dataType": "STRING",
        "isPartOfUniqueness": false
      },
      "value": ""
    },
    {
      "identifierType": {
        "name": "VM_LIMIT",
        "dataType": "STRING",
        "isPartOfUniqueness": false
      },
      "value": ""
    }
  ]
},
"description": "A vCenter Adapter Instance for VC 12.345.678.9",
"collectorId": 1,
"collectorGroupId": "909c2fbf-2c2c-4957-9a75-21bf2a887d31",
"credentialInstanceId": "65081a8d-d462-43b2-b4e0-596eaf3d497e",
"monitoringInterval": 5,
"adapter-certificates": [
  {
    "thumbprint": "2520fb4351bc91ee7b82ef7cc54a8d88fa893da9",
    "certificateDetails": "[
      Version: V3 Subject: C=US, CN=12.345.678.9
      Signature Algorithm: SHA256withRSA, OID = 1.2.840.113549.1.1.11
```

```
        Key: Sun RSA public key, 2048 bits modulus: ...
        Validity: [From: Wed Jul 15 19:26:51 UTC 2015, To: Tue Jul 08 11:26:30 UTC 2025]
        Issuer: O=W12R2UINanduVC, C=US, DC=local, DC=vsphere, CN=CA ...
          ...
      ]"
    }
  ],
  ...
  "id": "a97bd204-e3e5-404b-a219-e2b20cf158d2"
}
```

Response in JSON:

```
{
  "resourceKey": {
    "name": "VC Adapter Instance",
    "adapterKindKey": "VMWARE",
    "resourceKindKey": "VMwareAdapter Instance",
    "resourceIdentifiers": [
      {
        "identifierType": {
          "name": "AUTODISCOVERY",
          "dataType": "STRING",
          "isPartOfUniqueness": false
        },
        "value": "true"
      },
      {
        "identifierType": {
          "name": "DISABLE_COMPUTATION_BASED_ON_CONSUMERS",
          "dataType": "STRING",
          "isPartOfUniqueness": false
        },
        "value": ""
```

```
},
{
  "identifierType": {
    "name": "DV_PORT_GROUP_DISABLED",
    "dataType": "STRING",
    "isPartOfUniqueness": false
  },
  "value": ""
},
{
  "identifierType": {
  "name": "DVS_DISABLED",
  "dataType": "STRING",
  "isPartOfUniqueness": false
  },
  "value": ""
},
{
  "identifierType": {
    "name": "PROCESSCHANGEEVENTS",
    "dataType": "STRING",
    "isPartOfUniqueness": false
  },
  "value": "true"
},
{
  "identifierType": {
    "name": "VCURL",
    "dataType": "STRING",
    "isPartOfUniqueness": true
  },
  "value": "https://12.345.678.9"
```

```
      },
      {
        "identifierType": {
          "name": "VM_FOLDER_DISABLED",
          "dataType": "STRING",
          "isPartOfUniqueness": false
        },
        "value": ""
      },
      {
        "identifierType": {
          "name": "VM_LIMIT",
          "dataType": "STRING",
          "isPartOfUniqueness": false
        },
        "value": ""
      }
    ]
  },
  "description": "A vCenter Adapter Instance for VC 12.345.678.9",
  "collectorId": 1,
  "collectorGroupId": "909c2fbf-2c2c-4957-9a75-21bf2a887d31",
  "credentialInstanceId": "65081a8d-d462-43b2-b4e0-596eaf3d497e",
  "monitoringInterval": 5,
  ...
  "id": "a97bd204-e3e5-404b-a219-e2b20cf158d2"
}
```

The response is same as the request body without the `adapter-certificates` section.

## Start Monitoring the New Adapter Instance

After the creating the adapter instance and configuring VMware Aria OperationsVMware Cloud Foundation Operations to recognize a valid certificate, start monitoring and collecting data. Your PUT request provides the UUID of the adapter instance used to discover new objects.

Verify that you have the UUID of the newly created adapter instance. See the response in Adapter Instance.

1. Make a PUT request to start monitoring with the new adapter instance.

   ```
   PUT https://operations.example.com/suite-api/api/adapters/<adapter_UUID>/
   monitoringstate/start
   ```

   ```
   PUT https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapters/
   <adapter_UUID>/monitoringstate/start
   ```

## Discover Objects and Collect Data

This example starts the adapter monitoring process using the adapter instance ID from the PUT request that created the adapter instance.

Request header:

```
PUT https://operations.example.com/suite-api/api/adapters/a97bd204-e3e5-404b-a219-
e2b20cf158d2/monitoringstate/start
```

```
PUT https://www.mgmt.cloud.vmware.com/vrops-cloud/suite-api/api/adapters/a97bd204-
e3e5-404b-a219-e2b20cf158d2/monitoringstate/start
```

# Documentation Legal Notice

Information about the documentation legal notice.