vmware by Broadcom

VMware Cloud Director Container Service Extension 4.2

Table of Contents

Release Notes

Includes product enhancements and notices, bug fixes, and resolved issues.

VMware Cloud Director Container Service Extension 4.2.3 Release Notes

This document contains the following sections

Introduction

VMware Cloud Director Container Service Extension 4.2.3 | 09 OCT 2024 | Build 24322572
Check for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics.

What's New in October 2024

Kubernetes Container Clusters UI plug-in 4.2.3 is available

To upgrade the plug-in to version 4.2.3, perform the following tasks. For more information, see Managing Plug-Ins.

1. Log in to the Broadcom Support Portal and download the latest version of the Kubernetes Container Clusters UI plug-in.
2. In the VMware Cloud Director Portal, from the top navigation bar, select **More>Customize Portal**.
3. Select the check boxes next to earlier versions of the Kubernetes Container Clusters UI plug-in and click **Disable**.
4. Click **Upload** and in the **Upload Plugin** wizard, upload the latest Kubernetes Container Clusters UI plug-in zip file.
5. To start using the new plug-in, refresh your browser.

Added support for Tanzu Kubernetes Grid 2.5.2

The following Kubernetes versions are now supported.

- 1.26.14
- 1.27.15
- 1.28.11
- 1.29.6
- 1.30.2

Product Support Notice

VMware Cloud Director Container Service Extension Server 4.2.3 includes only the Kubernetes Container Clusters UI plug-in

This release does not include a new version of the VMware Cloud Director Container Service Extension Server OVA file.

Kubernetes Container Clusters UI plug-in 4.2.3 supports VMware Cloud Director 10.5 and 10.6 only

VMware Cloud Director Container Service Extension 4.2.2 supports VMware Cloud Director 10.5 and 10.6 only

To ensure compatibility with VMware Cloud Director 10.5 and 10.6, you must upgrade your existing version of VMware Cloud Director Container Service Extension to version 4.2.2.

The Auto Repair on Errors toggle is deprecated and is not supported in this release

Tanzu Kubernetes Grid versions 1.6.1, 1.5.4, and 1.4.3 are not supported

Tanzu Kubernetes Grid versions 1.6.1, 1.5.4 and 1.4.3 are no longer supported in VMware Cloud Director Container Service Extension 4.2 and later. For more information on the end of this support, see Product Lifecycle Matrix.

- New cluster deployments by using unsupported versions fail.
- Existing Tanzu Kubernetes Grid clusters must be upgraded by service providers or tenant users to version 2.1.1, 2.2, 2.3.1 or 2.4 and supported Kubernetes versions.

Compatibility Notices

VMware Cloud Director Container Service Extension 4.2.3 Interoperability Updates with Kubernetes Resources

To view the interoperability of VMware Cloud Director Container Service Extension 4.2.3 and previous versions with VMware Cloud Director, and additional product interoperability, see the Product Interoperability Matrix.

The following table displays the interoperability between VMware Cloud Director Container Service Extension 4.2.3 and Kubernetes resources.

Resources	Supported Versions	Documentation
Kubernetes Cloud Provider for VMware Cloud Director™	1.6.1	Kubernetes Cloud Provider for VMware Cloud Director Documentation
Kubernetes Container Storage Interface Driver for VMware Cloud Director™	1.6.0	Kubernetes Container Storage Interface driver for VMware Cloud Director
Kubernetes Cluster API Provider for VMware Cloud Director™	1.3.2	Kubernetes Cluster API Provider for VMware Cloud Director
RDE Projector	0.7.1	Not applicable

As a service provider, you can manually update Kubernetes resources.

1. In VMware Cloud Director UI, from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. In Kubernetes Container Clusters UI plug-in, select **CSE Management** > **Server Details** > **Update Server** > **Update Configuration** > **Next**.
3. In the **Current CSE Server Components** section, update the Kubernetes resources configuration.
4. Click **Submit Changes**.

For more information, see Update the VMware Cloud Director Container Service Extension Server documentation.

After you install or upgrade to VMware Cloud Director Container Service Extension 4.2.3 by using the Kubernetes Container Clusters UI plug-in, component versions of the VMware Cloud Director Container Service Extension server configuration are updated automatically

The following components versions are used in VMware Cloud Director Container Service Extension 4.2.3.

- kind: v0.24.0
- clusterctl: v1.7.4
- core capi: v1.7.4
- bootstrap provider: v1.7.4

- control plane provider: v1.7.4
- kindest: v1.31.0
- cert manager: v1.15.1

When attempting certain workflows in the Kubernetes Container Clusters UI plug-in, you might see a warning message

The following message might be displayed: `Confirm that the components in this cluster have the required versions.` You must verify that the relevant Kubernetes component versions listed on the **Cluster Information** page of the Kubernetes Container Clusters UI plug-in, match the supported versions in the above table.

- If the component versions match, ignore the message.
- If the component versions do not match, follow the instructions in Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters.

> **NOTE**
>
> For clusters that were created by using older versions of VMware Cloud Director Container Service Extension, perform a one time script upgrade action. This allows the clusters to be compatible with the latest VMware Cloud Director Container Service Extension.

VMware Cloud Director Container Service Extension 4.2.2 Release Notes

This document contains the following sections

Introduction

VMware Cloud Director Container Service Extension 4.2.2 | 27 JUN 2024 | Build 24053860 Check for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics.

What's New in June 2024

VMware Cloud Director Container Service Extension Server 4.2.2 is available

As a service provider, you can upgrade the VMware Cloud Director Container Service Extension Server to the latest version.

1. Log in to the Broadcom Support Portal and download VMware Cloud Director Container Service Extension Server 4.2.2 OVA file.
2. In the Kubernetes Container Clusters UI plug-in, select **CSE Management > Server Details > Update Server**.
3. Update your **VMware Cloud Director Container Service Extension Server.** See Patch Version Upgrade.

Kubernetes Container Clusters UI plug-in 4.2.2 is available

To upgrade the plug-in to version 4.2.2, perform the following tasks. For more information, see Managing Plug-Ins.

1. Log in to the Broadcom Support Portal and download the latest version of the Kubernetes Container Clusters UI plug-in.
2. In the VMware Cloud Director Portal, from the top navigation bar, select **More>Customize Portal**.
3. Select the check boxes next to earlier versions of the Kubernetes Container Clusters UI plug-in and click **Disable**.
4. Click **Upload** and in the **Upload Plugin** wizard, upload the latest Kubernetes Container Clusters UI plug-in zip file.
5. To start using the new plug-in, refresh your browser.

Product Support Notice

VMware Cloud Director Container Service Extension 4.2.2 supports VMware Cloud Director 10.5 and 10.6 only

To ensure compatibility with VMware Cloud Director 10.5 and 10.6, you must upgrade your existing version of VMware Cloud Director Container Service Extension to version 4.2.2.

Kubernetes Container Clusters UI plug-in 4.2.2 does not support VMware Cloud Director 10.4 and earlier

The Auto Repair on Errors toggle is deprecated and is not supported in this release

Tanzu Kubernetes Grid versions 1.6.1, 1.5.4, and 1.4.3 are not supported

Tanzu Kubernetes Grid versions 1.6.1, 1.5.4 and 1.4.3 are no longer supported by VMware in VMware Cloud Director Container Service Extension 4.2. For more information on the end of this support, see Product Lifecycle Matrix.

- New cluster deployments by using unsupported versions fail.
- Existing Tanzu Kubernetes Grid clusters must be upgraded by service providers or tenant users to version 2.1.1, 2.2, 2.3.1 or 2.4 and supported Kubernetes versions.

Compatibility Notices

VMware Cloud Director Container Service Extension 4.2.2 Interoperability Updates with Kubernetes Resources

To view the interoperability of VMware Cloud Director Container Service Extension 4.2.2 and previous versions with VMware Cloud Director, and additional product interoperability, see the Product Interoperability Matrix.

The following table displays the interoperability between VMware Cloud Director Container Service Extension 4.2.2 and Kubernetes resources.

Resources	Supported Versions	Documentation
Kubernetes Cloud Provider for VMware Cloud Director™	1.6.0	Kubernetes Cloud Provider for VMware Cloud Director Documentation
Kubernetes Container Storage Interface Driver for VMware Cloud Director™	1.6.0	Kubernetes Container Storage Interface driver for VMware Cloud Director
Kubernetes Cluster API Provider for VMware Cloud Director™	1.3.0	Kubernetes Cluster API Provider for VMware Cloud Director
RDE Projector	0.7.1	Not applicable

As a service provider, you can manually update Kubernetes resources.

1. In VMware Cloud Director UI, from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. In Kubernetes Container Clusters UI plug-in, select **CSE Management** > **Server Details** > **Update Server** > **Update Configuration** > **Next**.
3. In the **Current CSE Server Components** section, update the Kubernetes resources configuration.
4. Click **Submit Changes**.

For more information, see Update the VMware Cloud Director Container Service Extension Server documentation.

After you install or upgrade to VMware Cloud Director Container Service Extension 4.2.2 by using the Kubernetes Container Clusters UI plug-in, components of the VMware Cloud Director Container Service Extension server configuration are updated automatically

The following components versions are used in VMware Cloud Director Container Service Extension 4.2.2.

- kind: v0.20.0
- clusterctl: v1.5.4
- core capi: v1.5.4
- bootstrap provider: v1.5.4
- control plane provider: v1.5.4
- kindest: v1.27.3
- cert manager: v1.13.2

When attempting certain workflows in the Kubernetes Container Clusters UI plug-in, you might see a warning message

The following message might be displayed: `Confirm that the components in this cluster have the required versions.` You must verify that the relevant Kubernetes component versions listed on the **Cluster Information** page of the Kubernetes Container Clusters UI plug-in, match the supported versions in the above table.

- If the component versions match, ignore the message.
- If the component versions do not match, follow the instructions in Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters.

  NOTE

  For clusters that were created by using older versions of VMware Cloud Director Container Service Extension, perform a one time script upgrade action. This allows the clusters to be compatible with the latest VMware Cloud Director Container Service Extension.

Known Issues

When a Kubernetes cluster is connected to an Edge Gateway, that does not belong to data center group and supports IP spaces, force delete operation of the cluster might fail to release the IP allocations in the cluster

This issue occurs when using VMware Cloud Director Container Service Extension 4.2.2 and earlier with Kubernetes Cloud Provider for VMware Cloud Director 1.6.0 and earlier and Kubernetes Cluster API Provider for VMware Cloud Director 1.3.0 and earlier.

Workaround: In Kubernetes Container Clusters UI plug-in, use the delete operation. Alternatively, before force deleting a Kubernetes cluster, manually release the IP allocations.

In VMware Cloud Director Container Service Extension 4.2.2 and earlier, when a Kubernetes cluster is connected to a data center group organization VDC network, the delete and force delete operations of the cluster might fail

Workaround: None.

When using VMware Cloud Director Container Service Extension 4.2.2 with Kubernetes Cloud Provider for VMware Cloud Director 1.6.0 and Kubernetes Cluster API Provider for VMware Cloud Director 1.3.0, you might not be able to create a Kubernetes cluster connected to a data center group organization VDC network

Workaround: As a service provider, update Kubernetes Cloud Provider for VMware Cloud Director (CPI) to version 1.6.1 and Kubernetes Cluster API Provider for VMware Cloud Director (CAPVCD) to version 1.3.1.

1. In VMware Cloud Director UI, from the top navigation bar, select **More** > **Kubernetes Container Clusters**.

2. In Kubernetes Container Clusters UI plug-in, select **CSE Management** > **Server Details** > **Update Server** > **Update Configuration** > **Next**.
3. In the CSE Server Components section, update the your CAPVCD Version to `v1.3.1` and your Cloud Provider Interface (CPI) Version to `1.6.1`.
4. Click **Submit Changes**.

In VMware Cloud Director Container Service Extension Server 4.2.2, as a user in the system organization, you cannot create a Kubernetes cluster in a tenant organization

This issue occurs due to changes to tenant Runtime Defined Entity (RDE) creation by provider users. For more information, see VMware Cloud Director 10.6 release notes.

Workaround: As a system administrator, create a new user in the tenant organization and by using the tenant user, create a Kubernetes cluster.

In VMware Cloud Director Container Service Extension Server 4.2.2, you cannot delete a Kubernetes cluster that is owned by a user in the system organization

This issue occurs due to changes to tenant Runtime Defined Entity (RDE) creation by provider users. For more information, see VMware Cloud Director 10.6 release notes.

Workaround: Perform the following tasks.

1. In the Kubernetes Clusters UI plug-in, in the **Cluster Info** page of the affected cluster, note down the values of the following resources: `RDE ID`, `vApp ID`, `Virtual Services`, `Persistent Volumes`.
2. By using the VMware Cloud Director UI, manually delete the resources.
3. By using the VMware Cloud Director API, manually delete the RDE.

When upgrading to Tanzu Kubernetes Grid 2.5.0 and Kubernetes 1.26.11, the process might fail with an error message

Versions 4.2.0 and 4.2.1 of the Kubernetes Clusters UI plug-in use incorrect coreDNS versions for the following products.

- Tanzu Kubernetes Grid 2.4.0 and Kubernetes 1.25.13
- Tanzu Kubernetes Grid 2.4.0 and Kubernetes 1.26.8

As the coreDNS version of the target Kubernetes version is earlier than the source Kubernetes version, the Core CAPI component of the cluster restricts the upgrade. As a result, the control plane nodes are not upgraded and the overall process fails. For example, you might observe the following error message.

```
[admission  webhook }}\{{"validation.kubeadmcontrolplane.controlplane.cluster.x-
k8s.io"}}  \{{denied the request: KubeadmControlPlane.controlplane.cluster.x-
k8s.io  }}\{{"testtf4-control-plane-node-pool"}} \{{is invalid:
spec.kubeadmConfigSpec.clusterConfiguration.dns.imageTag: Forbidden:  cannot migrate
CoreDNS up to }}\{{'1.9.3'}} \{{from }}\{{{}'1.10.1'{}}}\{{{}: cannot migrate up
to  {}}}\{{'1.9.3'}} \{{from }}\{{{}'1.10.1'{}}}\{{{}] during patching objects with
name  [KubeadmControlPlane/testtf4-control-plane-node-pool]
```

Workaround: Perform one of the following tasks.

- If the upgrade process of your cluster failed, perform the following steps to resume the process.
  a. In the Kubernetes Clusters UI plug-in, in the **Cluster Info** page of the affected cluster, note down the `<clusterId>` value.
  b. In Postman, run `GET https://<vcd>/cloudapi/1.0.0/entities/<clusterId>` with ETag, where `<clusterId>` is the value from step 1 and `<vcd>` is the URL address of your VMware Cloud Director instance. For more information about use of ETags, see *Runtime Defined Entities and Behaviors* in the Cloud Director Extension SDK documentation.
  c. In the `capiYaml` contents of the RDE, at `entity.spec.capiYaml` JSON path, replace `v1.9.3_vmware.16` with `v1.10.1_vmware.13`.

   d. Run `PUT https://<vcd>/cloudapi/1.0.0/entities/<clusterId>` , where `<clusterId>` is the value from step 1 and `<vcd>` is the URL address of your VMware Cloud Director instance.
      Note: You must use the entire payload that you received as a result of the GET operation and the ETag from step 2 and the modified content from step 3.
- If performing a new upgrade, perform the upgrade by using the following path.
   a. From Tanzu Kubernetes Grid 2.4 and Kubernetes 1.25.13 upgrade to Tanzu Kubernetes Grid 2.4 and Kubernetes 1.26.8.
   b. From Tanzu Kubernetes Grid 2.4 and Kubernetes 1.26.8 upgrade to Tanzu Kubernetes Grid 2.5 and Kubernetes 1.27.8.

When using VMware Cloud Director Container Service Extension to create a cluster, the operation might fail

If you are using 172.17.0.0/16 and 172.18.0.0/16 CIDR ranges or IP addresses from these ranges in your external network pool, the creation of a cluster might fail after the first control plane VM is created and you might observe `guestinfo.cloudinit.target.cluster.get.kubeconfig.status` phase failures in the **Events** tab of the Kubernetes Container Clusters UI plug-in.

The ephemeral VMs that leverage Docker use the same CIDR ranges during the creation of the bootstrap cluster. As a result of the IP conflict, communication between the components of the bootstrap cluster and the control plane VM is affected, which causes the cluster creation to fail.

Workaround: Ensure that you are not using 172.17.0.0/16 and 172.18.0.0/16 CIDR ranges or IP addresses from these ranges in the following network assets.

- Organization VDC network ranges where your TKG clusters are deployed.
- External IP allocations and ranges that are used by the Organization Edge Gateway and the associated Load Balancer.
- Infrastructure networks where your DNS servers are connected.
- The IP address, which the VMware Cloud Director public API endpoint URL resolves to.

When you resize a disk volume by using online expansion in the Kubernetes Container Storage Interface Driver for VMware Cloud Director solution, the operation might fail

When attempting an online expansion of a volume on a named disk that is fast and thin provisioned, attached to a VM, and the name of the storage profile differs between the `StorageClass` and the VM, the `csi-resizer` container in the `csi-vcd-controllerplugin` pod might display the following error message. This is a known issue in VMware Cloud Director version 10.5.1.1 and earlier.

```
API Error: 400: [ ddedf59a-8efe-418f-9417-b4ce6aad2883 ] Cannot use multiple storage profiles in a fast-provisioned VDC "tenant_org_name" for VM "cluster-worker-node-pool-name".]
```

Workaround: Before performing online volume expansion, verify that the storage profile name for the VM are the same as specified in `StorageClass` .

After a cluster upgrade, the Kubernetes Container Storage Interface Driver for VMware Cloud Director solution does not run as expected

After running version v0.1.3 of the `cluster-upgrade-script-airgapped` container to upgrade a cluster, the images of the Kubernetes Container Storage Interface Driver for VMware Cloud Director solution are updated, but 1 of the 2 `nodeplugin` pods is in error state, such as `CrashLoopBackoff` or `Error` .

Workaround: To recreate the `nodeplugin` pods, replace and update the `DaemonSet` by running the following command.

```
kubectl replace --force -f "https://raw.githubusercontent.com/vmware/cloud-director-named-disk-csi-driver/1.6.0/manifests/csi-node.yaml"
```

In VMware Cloud Director Container Service Extension 4.2.1, if you force delete clusters that are configured to use IP Spaces, the IP allocated to the cluster and/or Kubernetes services running on the cluster are not released automatically, and manual intervention is necessary.

Workaround: None

The Create a Tanzu Kubernetes Grid Cluster and Create New Worker Node Pools workflows might fail

When both sizing policy and vGPU policy are specified and the vGPU policy already contains sizing information, the workflows cannot be completed successfully.

Workaround: If you select a vGPU policy that already contains sizing information during **Create a Tanzu Kubernetes Grid Cluster** workflow, or **Create New Worker Node Pools** workflow, do not also select a sizing policy.

VMware Cloud Director services fail continuously after startup

When a resolve operation is invoked on an RDE that has a lot of tasks associated with it, VMware Cloud Director crashes with the `java.lang.OutOfMemoryError: Java heap space` error message. The issue is present on VMware Cloud Director 10.4 and above. For more information, see VMware Knowledge Base Article 95464.

Workaround: None

Registry URL changes in VMware Cloud Director Container Service Extension configuration are not supported

Workaround: Use load balancers to front registry virtual machines to swap the virtual machines out if necessary.

If you use VMware Cloud Director 10.4.2.2, the cluster deletion workflow in Kubernetes Container Clusters UI plug-in might fail

The cluster deletion operation fails with the following error message.

```
"error": "failed to delete VCD Resource [clusterName] of type [VApp] from VCDResourceSet
of RDE [urn:vcloud:entity:vmware:capvcdCluster:<uuid>]: [failed to update capvcd status
for RDE [urn:vcloud:entity:vmware:capvcdCluster:<uuid>]; expected http response [200],
obtained [400]: resp: [\"{\\\"minorErrorCode\\\":\\\"BAD_REQUEST\\\",\\\"message\\\":\\
\"[ a8e89bd2-195d-458b-808d-3ff81e074fa0 ] RDE_CANNOT_VALIDATE_AGAINST_SCHEMA [ #/status/
capvcd/vcdResourceSet/2: expected type: JSONObject, found: Null\\\\n ]\\\",\\\"stackTrace\
\\":null}\"]: [400 Bad Request]]"
```

This is a known issue in VMware Cloud Director 10.4.2.2.

Workaround: Delete the cluster by using the **Force Delete** workflow.

In Kubernetes Container Clusters UI plug-in, the CSE Management upgrade workflows might add or remove rights from the CSE Admin Role or Kubernetes Cluster Author

If required rights are missing, users might face errors during cluster workflows.

Workaround: Manually update the **Custom** roles that are cloned from **CSE Admin Role** or **Kubernetes Cluster Author** role.

VMware Cloud Director Container Service Extension does not automatically install a Tanzu-standard repository in the Tanzu Kubernetes Grid 2.1.1 and 2.2 clusters

Workaround: Perform one of the following tasks.

- When using VMware Cloud Director Container Service Extension 4.1, manually install the repository and packages.
- Upgrade to VMware Cloud Director Container Service Extension 4.1.1a. In this version, clusters created with Tanzu Kubernetes Grid 2.1.1 and 2.2 automatically have the Tanzu-standard repository installed.

When tenants attempt certain actions with VMware Cloud Director Container Service Extension, the following error messages might be displayed

Warnings:

- `Cannot fetch provider configuration. Please contact your administrator.`
  Tenant users may see this warning, and be blocked when they try to create a cluster.
- `Node Health Check settings have not been configured by your provider.`
  Tenant users may see this warning when they try to activate **Node Health Check** during cluster creation or in the cluster settings.

These warnings can occur for the following reasons:

- The VMware Cloud Director Container Service Extension server has not finished starting up.
- The VMware Cloud Director Container Service Extension server has not yet published the server configuration to tenant organizations. The server configuration is published automatically every hour from the server startup as the server is running. Therefore, publishing to new tenant organizations that are created during hourly window occurs at the end of the hour.
- The tenant user's role does not have the following right: `View: VMWARE:VCDKECONFIG` . This right was added to the **Kubernetes Cluster Author** global role in VMware Cloud Director Container Service Extension 4.1.
- There was an unexpected error while fetching the server configuration.

Workaround: Perform the following tasks.

1. Ensure that the VMware Cloud Director Container Service Extension server is operating successfully.
2. Ensure the tenant user's role has the right `View: VMWARE:VCDKECONFIG` . Tenant users must log out of VMware Cloud Director, and log back in to activate any changes made to their role.
3. Wait for hourly publishing to new organizations.

In some instances, nodes cannot join clusters even when the cluster is in an available state

This issue can occur intermittently and the following error message appears in the **Events** tab of the cluster info page in Kubernetes Container Clusters UI.

`VcdMachineScriptExecutionError` with the following details: `script failed with status [x] and reason [Date Time 1 /root/node.sh: exit [x]]`

Workaround: For VMware Cloud Director Container Service Extension 4.1, there is a retry mechanism added that uses a retry feature from Cluster API which reduces the occurrence of this issue.

VMware Cloud Director Container Service Extension 4.1 does not support Dynamic Host Configuration Protocol (DHCP)

The cluster creation workflow in VMware Cloud Director Container Service Extension 4.1 fails if the cluster is connected to a routed organization VDC network that uses DHCP instead of static IP pool to distribute IPs to virtual machines.

Workaround: VMware Cloud Director Container Service Extension 4.1 only supports organization VDC networks in the following scenarios.

- If the VDC is routed.
- If the VDC uses static IP pool to distribute IPs to virtual machines that are connected to it.

It is not possible to activate GPU support in an air-gapped cluster

As VMware cannot redistribute nVidia packages, it is not possible to activate GPU support in an air-gapped cluster out of box. The failure occurs when the cluster attempts to download the nVidia binary from `nvidia.github.io` in the cloud initialization script.

Workaround: As a service provider, you can potentially consider allowing the cluster access to `nvidia.github.io` by using a proxy server.

`Audit_trail` table grows rapidly in the VMware Cloud Director database due to RDE modify events being too large

RDE modify events log the whole body of the RDE that has changed. These large events cause the `audit_trail` table to grow longer than necessary.

Workaround: Upgrade to VMware Cloud Director 10.3.3.4 or later and perform one of the following tasks.

- If you are using VMware Cloud Director 10.3.3.4, set the `audit.rde.diffOnly` configuration property to `True` .
- If you are using VMware Cloud Director 10.4.0 or later, no changes in the configuration properties are required.

VMware Cloud Director Container Service Extension 4.1 uses Kubernetes Cluster API Provider for VMware Cloud Director 1.1 and Kubernetes Cloud Provider for VMware Cloud Director 1.4 by default

Kubernetes Cluster API Provider for VMware Cloud Director 1.1 and Kubernetes Cloud Provider for VMware Cloud Director 1.4 do not support IP spaces.

Workaround: None

Tanzu Addons-Manager does not appear after upgrading to Tanzu Kubernetes Grid 2.2.0 with Kubernetes v1.24+

After you upgrade a VMware Cloud Director Container Service Extension 4.0.3 cluster from Tanzu Kubernetes Grid 1.6.1 with Kubernetes v1.23.x to Tanzu Kubernetes Grid 2.2 with Kubernetes v1.24.x, `tanzu-addon-controller-manager` pod is stuck at **PENDING** or `CrashLoopBackOff` state for the following reason:

```
Error from server (NotFound): packageinstalls.packaging.carvel.dev "addons-
manager.tanzu.vmware.com" not found
```

Workaround: Manually delete the `tanzu-addons-controller-manager` deployment and **PackageInstall** object.

1. Delete the deployment by running the following commands.
   ```
   kubectl get deployments -A kubectl delete deployment -n tkg-system   tanzu-addons-controller-manager
   ```
2. Delete the **PackageInstall** object by running the following commands.
   ```
   kubectl get packageinstall -A kubectl delete packageinstall -n tkg-system   tanzu-addons-manager
   ```

When a cluster creation process finished, the API request to delete the ephemeral VM might fail

An ephemeral VM is created during the cluster creation process and is deleted by VMware Cloud Director Container Service Extension, when the process is complete. VMware Cloud Director Container Service Extension re-attempts to delete the ephemeral VM for up to 15 minutes. If VMware Cloud Director Container Service Extension fails to delete the ephemeral VM after reattempting, the ephemeral VM remains in the cluster's vApp. In the **Events** tab of the cluster info page in the Kubernetes Container Clusters UI plug-in, the `EphemeralVMError` error message appears with the following details.

```
error deleting Ephemeral VM [EPHEMERAL-TEMP-VM] in vApp [cluster-vapp-name]: [reason for
failure]. The Epemeral VM needs to be cleaned up manually.
```

The reason for failure depends on the stage at which the ephemeral VM deletion failed.

Workaround: In the VMware Cloud Director UI, delete the ephemeral VM from the cluster's vApp.

1. Log in to the VMware Cloud Director Tenant Portal, and from VMware Cloud Director navigation menu, select **Data Centers**.
2. In the Virtual Data Center page, select the organization tile, and from the left navigation menu, select **vApps.**
3. In the vApps page, select the vApp of the cluster.
4. In the cluster information page, click the ellipse to the left of the Ephermal VM, and click **Delete**.

If the ephemeral VM is not manually cleaned up when a delete request is issued, the cluster delete operation fails. It is then necessary to force delete the cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster, and in the cluster information page, click **Delete**.
3. In the **Delete Cluster** page, select the **Force Delete** checkbox, and click **Delete**.

When using a direct organization VDC network with NSX in VMware Cloud Director, creating clusters in VMware Cloud Director Container Service Extension 4.1 is not possible

VMware Cloud Director Container Service Extension 4.1 clusters do not support this configuration.

Workaround: None

In VMware Cloud Director Container Service Extension, the creation of Tanzu Kubernetes Grid clusters can fail due to a script execution error

In the **Events** tab of the cluster info page in Kubernetes Container Clusters UI plug-in, the `ScriptExecutionTimeout` error message appears with the following details.

```
error while bootstrapping the machine [cluster-name/EPHEMERAL_TEMP_VM]; timeout for post customization phase [phase name of script execution]
```

Workaround: To re-attempt cluster creation, activate **Auto Repair on Errors** from cluster settings.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster, and in the cluster information page, click **Settings,** and activate the **Auto Repair on Errors** toggle.
3. Click **Save.**

   NOTE

   If you are troubleshooting issues related to cluster creation, deactivate the **Auto Repair on Errors** toggle.

In Kubernetes Container Clusters UI plug-in, when the cluster status is Error, the cluster delete operation might fail

Workaround: To delete a cluster in **Error** status, you must force delete the cluster.

1. Log in to VMware Cloud Director and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster and in the cluster information page, click **Delete**.
3. In the **Delete Cluster** page, select the **Force Delete** checkbox and click **Delete**.

Cluster creation fails with an error message

The **`ERROR: failed to create cluster: failed to pull image`** error message is displayed in the following scenarios.

- When a user attempts to create a Tanzu Kubernetes Grid Cluster using VMware Cloud Director Container Service Extension 4.1, and it fails intermittently.
- An image pull error due to a HTTP 408 response is reported.

This issue can occur if there is difficulty reaching the Internet from the EPHEMERAL_TEMP_VM to pull the required images.

Potential causes:

- Slow or intermittent Internet connectivity.
- The network IP Pool cannot resolve DNS (docker pull error).
- The network MTU behind a firewall must set lower.

Workaround: Ensure that there are no networking connectivity issues stopping the EPHEMERAL_TEMP_VM from reaching the Internet. For more information, refer to https://kb.vmware.com/s/article/90326.

Users may encounter authorization errors when executing cluster operations in Kubernetes Container Clusters UI plug-in if a Legacy Rights Bundle exists for their organization.

Workaround: Perform the following tasks.

1. After you upgrade VMware Cloud Director from version 9.1 or earlier, the system may create a **Legacy Rights Bundle** for each organization. This **Legacy Rights Bundle** includes the rights that are available in the associated organization at the time of the upgrade and is published only to this organization. To begin using the rights bundles model for an existing organization, you must delete the corresponding **Legacy Rights Bundle**.
2. In the **Administration** tab in the service provider portal, you can delete **Legacy Rights Bundles**. Kubernetes Container Clusters UI plug-in **CSE Management** has a server setup process that automatically creates, and publishes **Kubernetes Clusters Rights Bundle** to all tenants. The rights bundle contains all rights that are involved in Kubernetes cluster management in VMware Cloud Director Container Service Extension 4.0.

After selecting the purpose of policy modification, the policies selection in VMware Cloud Director Container Service Extension 4 plug-in does not populate the full list

When a user selects a sizing policy in the Kubernetes Container Clusters UI plug-in and they want to change it, the drop-down menu only displays the selected sizing policy, and does not automatically load alternative sizing policies. The user has to delete the text manually to allow the alternative sizing policies to appear. This also occurs in the drop-down menu when the user selects of placement policies and storage policies.

Workaround: None. This is intentional and typical behavior of the `combobox html` web component in Clarity, the web framework that VMware Cloud Director UI is built on. The drop-down box uses the input text as a filter. When the input field is empty, you can see all selections, and the selections filter as you type.

When you create a VMware Cloud Director Container Service Extension cluster, a character capitalization error appears

In the **Kubernetes Container Clusters** UI, if you use capital letters, the following error message appears.

```
Name must start with a letter, end with an alphanumeric, and only contain alphanumeric or hyphen (-) characters. (Max 63 characters)
```

Workaround: None. This is a restriction set by Kubernetes, where object names are validated under RFC 1035 labels. For more information, see the Kubernetes documentation.

Kubernetes Container Clusters UI plug-in 4.1 does not interoperate with other versions of the Kubernetes Container Clusters UI plug-in, such as 4.0 or 3.5.0

The ability to operate these two plug-ins simultaneously without conflict is a known limitation of the VMware Cloud Director UI. You can only have one plug-in activated at any given time.

Workaround: None.

VMware Cloud Director Container Service Extension fails to deploy clusters with TKG templates that have an unmodifiable placement policy set on them

Workaround: Perform the following tasks.

1. Log in to the VMware Cloud Director Tenant Portal as an administrator.
2. Click **Libraries** > **vApp Templates**.
3. In the **vApp Templates** window, select the radio button to the left of the template.
4. In the top ribbon, click **Tag with Compute Policies**.
5. Select the **Modifiable** check boxes, and click **Tag**.

In VMware Cloud Director 10.4, service providers cannot log in to the virtual machine of VMware Cloud Director Container Service Extension

In VMware Cloud Director 10.4, after deploying the VMware Cloud Director Container Service Extension virtual machine from OVA file, the following two check boxes in the VM settings page are not selected by default.

- Allow local administrator password
- Auto-generate password

Workaround: To allow service providers to log-in to the virtual machine of VMware Cloud Director Container Service Extension and perform troubleshooting tasks, select the Allow local administrator password and Auto-generate password check boxes.

1. Log in to VMware Cloud Director UI as a service provider and create a vApp from the VMware Cloud Director Container Service Extension OVA file.
2. After you deploy the vApp, and before you power it on, browse to **VM details** > **Guest OS Customization** and select **Allow local administrator password** and **Auto-generate password**.
3. After the update task finishes, power on the vApp.

Fast provisioned disks in Organization VCD cannot be resized

Workaround: To resize disks, deactivate fast provisioning in Organization VDC.

1. Log in to VMware Cloud Director UI as a provider, and select **Resources.**
2. In the **Cloud Resources** tab, select **Organization VDCs**, and select an organization VDC.
3. In the organization VDC window, under **Policies**, select **Storage**.
4. Click **Edit**, and deactivate the **Fast provisioning** toggle.
5. Click **Save**.

After you log in as a service provider and upload the latest Kubernetes Container Clusters UI plug-in, the CSE Management tab is not displayed

If there are multiple activated Kubernetes Container Clusters UI plug-ins with the same name or id but different versions, the lowest version of the plug-in is used. Only the highest version of the Kubernetes Container Clusters UI plug-in must be active. For more information on managing plug-ins, see Managing Plug-Ins.

Workaround: Deactivate the previous Kubernetes Container Clusters UI plug-ins.

1. Log in to **VMware Cloud Director** UI as a provider, and select **More > Customize Portal.**
2. Select the check box next to the names of the target plug-ins, and click **Enable** or **Disable**.
3. To start using the newly activated plug-in, refresh the Internet browser page.

Resize or upgrade a Tanzu Kubernetes Grid cluster by using kubectl

After a cluster is created in the Kubernetes Container Clusters UI plug-in, you can resize, upgrade, lifecycle manage the cluster, or manage workloads, by using `kubectl` instead of the Kubernetes Container Clusters UI plug-in.

1. To delete the RDE-Projector operator from the cluster, run `kubectl delete deployment -n rdeprojector-system rdeprojector-controller-manager`.
2. Detach the Tanzu Kubernetes Grid cluster from Kubernetes Container Clusters UI plug-in.
   a. In the VMware Cloud Director UI, in the **Cluster Overview** page, retrieve the cluster ID of the cluster.
   b. Update the RDE and set the `entity.spec.vcdKe.isVCDKECluster` value to false.
      a. To get the payload of the cluster, run `GET https://<vcd>/cloudapi/1.0.0/entities/<Cluster ID>`.
      b. Copy and update the json path in the payload.
      c. Set the `entity.spec.vcdKe.isVCDKECluster` value to false.
      d. Run `PUT https://<vcd>/cloudapi/1.0.0/entities/<Cluster ID>` with the modified payload. It is necessary to include the entire payload as the body of PUT operation.

      NOTE

      After performing the tasks above, the cluster is detached from VMware Cloud Director Container Service Extension 4.1 and you cannot manage the cluster through VMware Cloud Director Container Service Extension 4.1. You must use kubectl to manage, resize or upgrade the cluster by directly applying the cluster API specification, CAPI yaml.

VMware Cloud Director Container Service Extension 4.2.1 Release Notes

This document contains the following sections

Introduction

VMware Cloud Director Container Service Extension 4.2.1 | 21 MAR 2024 | Build 23512589
Check for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics.

What's New in March 2024

- **New version of the `getting-started-airgapped` container**: You can use version v0.1.3 of the `getting-started-airgapped` container to configure the private registry for air-gapped functionality. For more information, see Create an Airgapped Environment.
- **New version of the `cluster-upgrade-script-airgapped` container**: You can use version v0.1.3 of the `cluster-upgrade-script-airgapped` container to update the Kubernetes components in the pre-existing cluster to the recommended versions. For more information, see Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters.
- **Additional support for Tanzu Kubernetes Grid and Kubernetes**: You can use Tanzu Kubernetes Grid 2.5 with Kubernetes versions 1.26, 1.27, and 1.28.
- **New VMware Cloud Director IP spaces feature**: As a service provider you must add the rights to use the feature to the **Kubernetes Cluster Author** role and the **Kubernetes Clusters Rights Bundle**. For more information, see Kubernetes Cluster Author Role and Kubernetes Clusters Rights Bundle. Additionally, you must use Kubernetes Cloud Provider for VMware Cloud Director 1.6 and Kubernetes Cluster API Provider for VMware Cloud Director v1.3.0.
- **Support for Online and Offline Volume Expansion**: With Kubernetes Container Storage Interface Driver for VMware Cloud Director 1.6.0, you can resize disks by performing online and offline volume expansion. To activate this feature, by using the **Kubernetes Cluster Author** role, set the `allowVolumeExpansion` field in `storageClass` to `true` . For more information, see the Kubernetes documentation.

Product Support Notice

VMware Cloud Director 10.4 and earlier are not supported with Kubernetes Container Clusters UI plug-in 4.2.2

The **Auto Repair on Errors** toggle is deprecated and will not be supported in the next version of the VMware Cloud Director Container Service Extension

Tanzu Kubernetes Grid versions 1.6.1, 1.5.4, and 1.4.3 are not supported

Tanzu Kubernetes Grid versions 1.6.1, 1.5.4 and 1.4.3 are no longer supported by VMware in VMware Cloud Director Container Service Extension 4.2. For more information on the end of this support, see Product Lifecycle Matrix.

- New cluster deployments by using unsupported versions fail.
- Existing Tanzu Kubernetes Grid clusters must be upgraded by service providers or tenant users to version 2.1.1, 2.2, 2.3.1 or 2.4 and supported Kubernetes versions.

Upgrade Notices

Kubernetes Container Clusters UI plug-in 4.2.1 is available to use with VMware Cloud Director

Before upgrading your VMware Cloud Director Container Service Extension server, you must first upgrade the Kubernetes Container Clusters UI plug-in. To upgrade the plug-in from version 4.2 to 4.2.1, perform the following tasks. For more information, see Managing Plug-Ins.

1. Log in to the Broadcom Support Portal and download version 4.2.1 of the Kubernetes Container Clusters UI plug-in.
2. In the VMware Cloud Director Portal, from the top navigation bar, select **More>Customize Portal**.
3. Select the check box next to Kubernetes Container Clusters UI plug-in 4.2 and click **Disable**.
4. Click **Upload** and in the **Upload Plugin** wizard, upload the Kubernetes Container Clusters UI plug-in 4.2.1 zip file.
5. To start using the new plug-in, refresh your browser.

VMware Cloud Director Container Service Extension Server 4.2.1 is available

As a service provider, you can upgrade the VMware Cloud Director Container Service Extension Server to version 4.2.1.

1. Log in to the Broadcom Support Portal and download VMware Cloud Director Container Service Extension Server 4.2.1.
2. In the Kubernetes Container Clusters UI plug-in of VMware Cloud Director, select **CSE Management > Server Details > Update Server**.
3. Update your **VMware Cloud Director Container Service Extension Server.**
   - To upgrade from 4.1.1a to 4.2.1, see Minor Version Upgrade.
   - To upgrade from 4.2 to 4.2.1, see Patch Version Upgrade.

Compatibility Notices

VMware Cloud Director Container Service Extension 4.2.1 Interoperability Updates with Kubernetes Resources

To view the interoperability of VMware Cloud Director Container Service Extension 4.2.1 and previous versions with VMware Cloud Director, and additional product interoperability, see the Product Interoperability Matrix.

The following table displays the interoperability between VMware Cloud Director Container Service Extension 4.2.1 and Kubernetes resources.

Kubernetes Resources	Supported Versions	Documentation
Kubernetes Cloud Provider for VMware Cloud Director™	1.6.0	Kubernetes Cloud Provider for VMware Cloud Director Documentation
Kubernetes Container Storage Interface Driver for VMware Cloud Director™	1.6.0	https://github.com/vmware/cloud-director-named-disk-csi-driver#container-storage-interface-csi-driver-for-vmware-cloud-director-named-independent-disks
Kubernetes Cluster API Provider for VMware Cloud Director™	v1.3.0	https://github.com/vmware/cluster-api-provider-cloud-director
RDE Projector	0.7.0	Not applicable

As a service provider, you can manually update Kubernetes resources by performing the following tasks.

1. In VMware Cloud Director UI, from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. In Kubernetes Container Clusters UI plug-in 4.2.1, select **CSE Management** > **Server Details** > **Update Server** > **Update Configuration** > **Next**.
3. In the **Current CSE Server Components** section, update the Kubernetes resources configuration.
4. Click **Submit Changes**.

For more information, see Update the VMware Cloud Director Container Service Extension Server documentation.

After you install or upgrade to VMware Cloud Director Container Service Extension 4.2.1 by using the Kubernetes Container Clusters UI plug-in, components of the VMware Cloud Director Container Service Extension server configuration are updated automatically

The following components versions are used in VMware Cloud Director Container Service Extension 4.2.1.

- kind: v0.20.0
- clusterctl: v1.5.4
- core capi: v1.5.4
- bootstrap provider: v1.5.4
- control plane provider: v1.5.4
- kindest: v1.27.3
- cert manager: v1.13.2

When attempting certain workflows in the Kubernetes Container Clusters UI plug-in, you might see a warning message

When tenant users attempt certain workflows in the Kubernetes Container Clusters UI plug-in, the following message might be displayed: `Confirm that the components in this cluster have the required versions.` You must verify that the relevant Kubernetes component versions listed on the **Cluster Information** page of the Kubernetes Container Clusters UI plug-in, match the supported versions in the above table.

- If the component versions match, ignore the message.
- If the component versions do not match, follow the instructions in Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters.

  NOTE

  For clusters that were created by using older versions of VMware Cloud Director Container Service Extension, perform a one time script upgrade action. This allows the clusters to be compatible with the latest VMware Cloud Director Container Service Extension.

Resolved Issues

Resolved Issues in 4.2.1

After creating a cluster with Node Health Check activated or activating Node Health Check for an existing cluster, the cluster cannot be managed

The Kubernetes Container Clusters UI plug-in 4.2 creates the `MachineHealthCheck` capi yaml object with an invalid `apiVersion` value. Clusters that contain invalid `MachineHealthCheck` section cannot be managed. This issue is fixed in the VMware Cloud Director Container Service Extension 4.2.1 release. You can create new clusters with **Node Health Check** activated and for existing clusters affected by this issue, to fix the value of the `MachineHealthCheck` capi yaml object, you can toggle **Node Health Check** on or off.

`Error` value is displayed on the cluster list page and on the Cluster Information page

When you select a cluster in the cluster list datagrid, you might see and `Error` in the **Upgrade** column. Additionally, on the **Cluster Information** page, the **Upgrade Availability** property displays `Error` . This issue is fixed in the VMware Cloud Director Container Service Extension 4.2.1 release. vApp templates that are invalid or corrupted are not acknowledged, and the upgrade availability value does not error when encountering such vApp templates.

Kubernetes templates datagrid fails to load with an error message

If the Kubernetes templates datagrid fails to load, the following error message is displayed: `Error: Failed to fetch Kubernetes Templates` . This issue is fixed in the VMware Cloud Director Container Service Extension 4.2.1 release.

vApp templates that are invalid or corrupted are not acknowledged, and the Kubernetes templates datagrid does not error when encountering such vapp templates.

Known Issues

When upgrading to Tanzu Kubernetes Grid 2.5.0 and Kubernetes 1.26.11, the process might fail with an error message

Versions 4.2.0 and 4.2.1 of the Kubernetes Clusters plug-in use incorrect coreDNS versions for the following products.

- Tanzu Kubernetes Grid 2.4.0 and Kubernetes 1.25.13
- Tanzu Kubernetes Grid 2.4.0 and Kubernetes 1.26.8

As the coreDNS version of the target Kubernetes version is earlier than the source Kubernetes version, the Core CAPI component of the cluster restricts the upgrade. As a result, the control plane nodes are not upgraded and the overall process fails. For example, you might observe the following error message.

```
[admission  webhook }}\{{"validation.kubeadmcontrolplane.controlplane.cluster.x-
k8s.io"}}  \{{denied the request: KubeadmControlPlane.controlplane.cluster.x-
k8s.io  }}\{{"testtf4-control-plane-node-pool"}} \{{is invalid:
spec.kubeadmConfigSpec.clusterConfiguration.dns.imageTag: Forbidden:  cannot migrate
CoreDNS up to }}\{{'1.9.3'}} \{{from }}\{{{}'1.10.1'{}}}\{{{}: cannot migrate up
to  {}}}\{{'1.9.3'}} \{{from }}\{{{}'1.10.1'{}}}\{{{}] during patching objects with
name  [KubeadmControlPlane/testtf4-control-plane-node-pool]
```

Workaround: Perform one of the following tasks.

- If the upgrade process of your cluster failed, perform the following steps to resume the process.
  a. In the Kubernetes Clusters plug-in, in the **Cluster Info** page of the affected cluster, note down the `<clusterId>` value.
  b. In Postman, run `GET https://<vcd>/cloudapi/1.0.0/entities/<clusterId>` with ETag, where `<clusterId>` is the value from step 1 and `<vcd>` is the URL address of your VMware Cloud Director instance. For more information about use of ETags, see *Runtime Defined Entities and Behaviors* in the Cloud Director Extension SDK documentation.
  c. In the `capiYaml` contents of the RDE, at `entity.spec.capiYaml` JSON path, replace `v1.9.3_vmware.16` with `v1.10.1_vmware.13`.
  d. Run `PUT https://<vcd>/cloudapi/1.0.0/entities/<clusterId>`, where `<clusterId>` is the value from step 1 and `<vcd>` is the URL address of your VMware Cloud Director instance.
  Note: You must use the entire payload that you received as a result of the GET operation and the ETag from step 2 and the modified content from step 3.
- If performing a new upgrade, perform the upgrade by using the following path.
  a. From Tanzu Kubernetes Grid 2.4 and Kubernetes 1.25.13 upgrade to Tanzu Kubernetes Grid 2.4 and Kubernetes 1.26.8.
  b. From Tanzu Kubernetes Grid 2.4 and Kubernetes 1.26.8 upgrade to Tanzu Kubernetes Grid 2.5 and Kubernetes 1.27.8.

When using VMware Cloud Director Container Service Extension to create a cluster, the operation might fail

If you are using 172.17.0.0/16 and 172.18.0.0/16 CIDR ranges or IP addresses from these ranges in your external network pool, the creation of a cluster might fail after the first control plane VM is created and you might observe `guestinfo.cloudinit.target.cluster.get.kubeconfig.status` phase failures in the **Events** tab of the Kubernetes Container Clusters UI plug-in.

The ephemeral VMs that leverage Docker use the same CIDR ranges during the creation of the bootstrap cluster. As a result of the IP conflict, communication between the components of the bootstrap cluster and the control plane VM is affected, which causes the cluster creation to fail.

Workaround: Ensure that you are not using 172.17.0.0/16 and 172.18.0.0/16 CIDR ranges or IP addresses from these ranges in the following network assets.

- Organization VDC network ranges where your TKG clusters are deployed.
- External IP allocations and ranges that are used by the Organization Edge Gateway and the associated Load Balancer.
- Infrastructure networks where your DNS servers are connected.
- The IP address, which the VMware Cloud Director public API endpoint URL resolves to.

When you resize a disk volume by using online expansion in the Kubernetes Container Storage Interface Driver for VMware Cloud Director solution, the operation might fail

When attempting an online expansion of a volume on a named disk that is fast and thin provisioned, attached to a VM, and the name of the storage profile differs between the `StorageClass` and the VM, the `csi-resizer` container in the `csi-vcd-controllerplugin` pod might display the following error message. This is a known issue in VMware Cloud Director version 10.5.1.1 and earlier.

```
API Error: 400: [ ddedf59a-8efe-418f-9417-b4ce6aad2883 ] Cannot use multiple storage profiles in a fast-provisioned VDC "tenant_org_name" for VM "cluster-worker-node-pool-name".]
```

Workaround: Before performing online volume expansion, verify that the storage profile name for the VM are the same as specified in `StorageClass` .

After a cluster upgrade, the Kubernetes Container Storage Interface Driver for VMware Cloud Director solution does not run as expected

After running version v0.1.3 of the `cluster-upgrade-script-airgapped` container to upgrade a cluster, the images of the Kubernetes Container Storage Interface Driver for VMware Cloud Director solution are updated, but 1 of the 2 `nodeplugin` pods is in error state, such as `CrashLoopBackoff` or `Error` .

Workaround: To recreate the `nodeplugin` pods, replace and update the `DaemonSet` by running the following command.

```
kubectl replace --force -f "https://raw.githubusercontent.com/vmware/cloud-director-named-disk-csi-driver/1.6.0/manifests/csi-node.yaml"
```

In VMware Cloud Director Container Service Extension 4.2.1, if you force delete clusters that are configured to use IP Spaces, the IP allocated to the cluster and/or Kubernetes services running on the cluster are not released automatically, and manual intervention is necessary.

Workaround: None

The Create a Tanzu Kubernetes Grid Cluster and Create New Worker Node Pools workflows might fail

When both sizing policy and vGPU policy are specified and the vGPU policy already contains sizing information, the workflows cannot be completed successfully.

Workaround: If you select a vGPU policy that already contains sizing information during **Create a Tanzu Kubernetes Grid Cluster** workflow, or **Create New Worker Node Pools** workflow, do not also select a sizing policy.

VMware Cloud Director services fail continuously after startup

When a resolve operation is invoked on an RDE that has a lot of tasks associated with it, VMware Cloud Director crashes with the `java.lang.OutOfMemoryError: Java heap space` error message. The issue is present on VMware Cloud Director 10.4 and above. For more information, see VMware Knowledge Base Article 95464.

Workaround: None

Registry URL changes in VMware Cloud Director Container Service Extension configuration are not supported

Workaround: Use load balancers to front registry virtual machines to swap the virtual machines out if necessary.

If you use VMware Cloud Director 10.4.2.2, the cluster deletion workflow in Kubernetes Container Clusters UI plug-in might fail

The cluster deletion operation fails with the following error message.

```
"error": "failed to delete VCD Resource [clusterName] of type [VApp] from VCDResourceSet
of RDE [urn:vcloud:entity:vmware:capvcdCluster:<uuid>]: [failed to update capvcd status
for RDE [urn:vcloud:entity:vmware:capvcdCluster:<uuid>]; expected http response [200],
obtained [400]: resp: [\"{\\\"minorErrorCode\\\":\\\"BAD_REQUEST\\\",\\\"message\\\":\\
\"[ a8e89bd2-195d-458b-808d-3ff81e074fa0 ] RDE_CANNOT_VALIDATE_AGAINST_SCHEMA [ #/status/
capvcd/vcdResourceSet/2: expected type: JSONObject, found: Null\\\\n ]\\\",\\\"stackTrace\
\\":null}\"]: [400 Bad Request]]"
```

This is a known issue in VMware Cloud Director 10.4.2.2. For more information, see VMware Cloud Director 10.4.2.2 Known Issues.

Workaround: Delete the cluster by using the **Force Delete** workflow.

In Kubernetes Container Clusters UI plug-in, the CSE Management upgrade workflows might add or remove rights from the CSE Admin Role or Kubernetes Cluster Author

If required rights are missing, users might face errors during cluster workflows.

Workaround: Manually update the **Custom** roles that are cloned from **CSE Admin Role** or **Kubernetes Cluster Author** role.

VMware Cloud Director Container Service Extension does not automatically install a Tanzu-standard repository in the Tanzu Kubernetes Grid 2.1.1 and 2.2 clusters

Workaround: Perform one of the following tasks.

- When using VMware Cloud Director Container Service Extension 4.1, manually install the repository and packages.
- Upgrade to VMware Cloud Director Container Service Extension 4.1.1a. In this version, clusters created with Tanzu Kubernetes Grid 2.1.1 and 2.2 automatically have the Tanzu-standard repository installed.

When tenants attempt certain actions with VMware Cloud Director Container Service Extension, the following error messages might be displayed

Warnings:

- `Cannot fetch provider configuration. Please contact your administrator.`
  Tenant users may see this warning, and be blocked when they try to create a cluster.
- `Node Health Check settings have not been configured by your provider.`
  Tenant users may see this warning when they try to activate **Node Health Check** during cluster creation or in the cluster settings.

These warnings can occur for the following reasons:

- The VMware Cloud Director Container Service Extension server has not finished starting up.
- The VMware Cloud Director Container Service Extension server has not yet published the server configuration to tenant organizations. The server configuration is published automatically every hour from the server startup as the server is running. Therefore, publishing to new tenant organizations that are created during hourly window occurs at the end of the hour.
- The tenant user's role does not have the following right: `View: VMWARE:VCDKECONFIG` . This right was added to the **Kubernetes Cluster Author** global role in VMware Cloud Director Container Service Extension 4.1.
- There was an unexpected error while fetching the server configuration.

Workaround: Perform the following tasks.

1. Ensure that the VMware Cloud Director Container Service Extension server is operating successfully.
2. Ensure the tenant user's role has the right `View: VMWARE:VCDKECONFIG`. Tenant users must log out of VMware Cloud Director, and log back in to activate any changes made to their role.
3. Wait for hourly publishing to new organizations.

In some instances, nodes cannot join clusters even when the cluster is in an available state

This issue can occur intermittently and the following error message appears in the **Events** tab of the cluster info page in Kubernetes Container Clusters UI.

`VcdMachineScriptExecutionError` with the following details: `script failed with status [x] and reason [Date Time 1 /root/node.sh: exit [x]]`

Workaround: For VMware Cloud Director Container Service Extension 4.1, there is a retry mechanism added that uses a retry feature from Cluster API which reduces the occurrence of this issue.

VMware Cloud Director Container Service Extension 4.1 does not support Dynamic Host Configuration Protocol (DHCP)

The cluster creation workflow in VMware Cloud Director Container Service Extension 4.1 fails if the cluster is connected to a routed organization VDC network that uses DHCP instead of static IP pool to distribute IPs to virtual machines.

Workaround: VMware Cloud Director Container Service Extension 4.1 only supports organization VDC networks in the following scenarios.

- If the VDC is routed.
- If the VDC uses static IP pool to distribute IPs to virtual machines that are connected to it.

It is not possible to activate GPU support in an air-gapped cluster

As VMware cannot redistribute nVidia packages, it is not possible to activate GPU support in an air-gapped cluster out of box. The failure occurs when the cluster attempts to download the nVidia binary from `nvidia.github.io` in the cloud initialization script.

Workaround: As a service provider, you can potentially consider allowing the cluster access to `nvidia.github.io` by using a proxy server.

`Audit_trail` table grows rapidly in the VMware Cloud Director database due to RDE modify events being too large

RDE modify events log the whole body of the RDE that has changed. These large events cause the `audit_trail` table to grow longer than necessary.

Workaround: Upgrade to VMware Cloud Director 10.3.3.4 or later and perform one of the following tasks.

- If you are using VMware Cloud Director 10.3.3.4, set the `audit.rde.diffOnly` configuration property to `True`.
- If you are using VMware Cloud Director 10.4.0 or later, no changes in the configuration properties are required.

VMware Cloud Director Container Service Extension 4.1 uses Kubernetes Cluster API Provider for VMware Cloud Director 1.1 and Kubernetes Cloud Provider for VMware Cloud Director 1.4 by default

Kubernetes Cluster API Provider for VMware Cloud Director 1.1 and Kubernetes Cloud Provider for VMware Cloud Director 1.4 do not support IP spaces.

Workaround: None

Tanzu Addons-Manager does not appear after upgrading to Tanzu Kubernetes Grid 2.2.0 with Kubernetes v1.24+

After you upgrade a VMware Cloud Director Container Service Extension 4.0.3 cluster from Tanzu Kubernetes Grid 1.6.1 with Kubernetes v1.23.x to Tanzu Kubernetes Grid 2.2 with Kubernetes v1.24.x, `tanzu-addon-controller-manager` pod is stuck at **PENDING** or `CrashLoopBackOff` state for the following reason:

```
Error from server (NotFound): packageinstalls.packaging.carvel.dev "addons-
manager.tanzu.vmware.com" not found
```

Workaround: Manually delete the `tanzu-addons-controller-manager` deployment and **PackageInstall** object.

1. Delete the deployment by running the following commands.
   ```
   kubectl get deployments -A kubectl delete deployment -n tkg-system   tanzu-addons-controller-manager
   ```
2. Delete the **PackageInstall** object by running the following commands.
   ```
   kubectl get packageinstall -A kubectl delete packageinstall -n tkg-system   tanzu-addons-manager
   ```

When a cluster creation process finished, the API request to delete the ephemeral VM might fail

An ephemeral VM is created during the cluster creation process and is deleted by VMware Cloud Director Container Service Extension, when the process is complete. VMware Cloud Director Container Service Extension re-attempts to delete the ephemeral VM for up to 15 minutes. If VMware Cloud Director Container Service Extension fails to delete the ephemeral VM after reattempting, the ephemeral VM remains in the cluster's vApp. In the **Events** tab of the cluster info page in the Kubernetes Container Clusters UI plug-in, the `EphemeralVMError` error message appears with the following details.

```
error deleting Ephemeral VM [EPHEMERAL-TEMP-VM] in vApp [cluster-vapp-name]: [reason for
failure]. The Epemeral VM needs to be cleaned up manually.
```

The reason for failure depends on the stage at which the ephemeral VM deletion failed.

Workaround: In the VMware Cloud Director UI, delete the ephemeral VM from the cluster's vApp.

1. Log in to the VMware Cloud Director Tenant Portal, and from VMware Cloud Director navigation menu, select **Data Centers**.
2. In the Virtual Data Center page, select the organization tile, and from the left navigation menu, select **vApps.**
3. In the vApps page, select the vApp of the cluster.
4. In the cluster information page, click the ellipse to the left of the Ephermal VM, and click **Delete**.

If the ephemeral VM is not manually cleaned up when a delete request is issued, the cluster delete operation fails. It is then necessary to force delete the cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster, and in the cluster information page, click **Delete**.
3. In the **Delete Cluster** page, select the **Force Delete** checkbox, and click **Delete**.

When using a direct organization VDC network with NSX in VMware Cloud Director, creating clusters in VMware Cloud Director Container Service Extension 4.1 is not possible

VMware Cloud Director Container Service Extension 4.1 clusters do not support this configuration.

Workaround: None

In VMware Cloud Director Container Service Extension, the creation of Tanzu Kubernetes Grid clusters can fail due to a script execution error

In the **Events** tab of the cluster info page in Kubernetes Container Clusters UI plug-in, the `ScriptExecutionTimeout` error message appears with the following details.

```
error while bootstrapping the machine [cluster-name/EPHEMERAL_TEMP_VM]; timeout for post
customization phase [phase name of script execution]
```

Workaround: To re-attempt cluster creation, activate **Auto Repair on Errors** from cluster settings.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster, and in the cluster information page, click **Settings,** and activate the **Auto Repair on Errors** toggle.
3. Click **Save.**

   NOTE

   If you are troubleshooting issues related to cluster creation, deactivate the **Auto Repair on Errors** toggle.

In Kubernetes Container Clusters UI plug-in, when the cluster status is Error, the cluster delete operation might fail

Workaround: To delete a cluster in **Error** status, you must force delete the cluster.

1. Log in to VMware Cloud Director and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster and in the cluster information page, click **Delete**.
3. In the **Delete Cluster** page, select the **Force Delete** checkbox and click **Delete**.

Cluster creation fails with an error message

The **`ERROR: failed to create cluster: failed to pull image`** error message is displayed in the following scenarios.

- When a user attempts to create a Tanzu Kubernetes Grid Cluster using VMware Cloud Director Container Service Extension 4.1, and it fails intermittently.
- An image pull error due to a HTTP 408 response is reported.

This issue can occur if there is difficulty reaching the Internet from the EPHEMERAL_TEMP_VM to pull the required images.

Potential causes:

- Slow or intermittent Internet connectivity.
- The network IP Pool cannot resolve DNS (docker pull error).
- The network MTU behind a firewall must set lower.

Workaround: Ensure that there are no networking connectivity issues stopping the EPHEMERAL_TEMP_VM from reaching the Internet. For more information, refer to https://kb.vmware.com/s/article/90326.

Users may encounter authorization errors when executing cluster operations in Kubernetes Container Clusters UI plug-in if a Legacy Rights Bundle exists for their organization.

Workaround: Perform the following tasks.

1. After you upgrade VMware Cloud Director from version 9.1 or earlier, the system may create a **Legacy Rights Bundle** for each organization. This **Legacy Rights Bundle** includes the rights that are available in the associated organization at the time of the upgrade and is published only to this organization. To begin using the rights bundles model for an existing organization, you must delete the corresponding **Legacy Rights Bundle**.
2. In the **Administration** tab in the service provider portal, you can delete **Legacy Rights Bundles**. Kubernetes Container Clusters UI plug-in **CSE Management** has a server setup process that automatically creates, and publishes **Kubernetes Clusters Rights Bundle** to all tenants. The rights bundle contains all rights that are involved in Kubernetes cluster management in VMware Cloud Director Container Service Extension 4.0.

After selecting the purpose of policy modification, the policies selection in VMware Cloud Director Container Service Extension 4 plug-in does not populate the full list

When a user selects a sizing policy in the Kubernetes Container Clusters UI plug-in and they want to change it, the drop-down menu only displays the selected sizing policy, and does not automatically load alternative sizing policies. The user

has to delete the text manually to allow the alternative sizing policies to appear. This also occurs in the drop-down menu when the user selects of placement policies and storage policies.

Workaround: None. This is intentional and typical behavior of the `combobox html` web component in Clarity, the web framework that VMware Cloud Director UI is built on. The drop-down box uses the input text as a filter. When the input field is empty, you can see all selections, and the selections filter as you type.

When you create a VMware Cloud Director Container Service Extension cluster, a character capitalization error appears

In the **Kubernetes Container Clusters** UI, if you use capital letters, the following error message appears.

```
Name must start with a letter, end with an alphanumeric, and only contain alphanumeric or hyphen (-) characters. (Max 63 characters)
```

Workaround: None. This is a restriction set by Kubernetes, where object names are validated under RFC 1035 labels. For more information, see the Kubernetes documentation.

Kubernetes Container Clusters UI plug-in 4.1 does not interoperate with other versions of the Kubernetes Container Clusters UI plug-in, such as 4.0 or 3.5.0

The ability to operate these two plug-ins simultaneously without conflict is a known limitation of the VMware Cloud Director UI. You can only have one plug-in activated at any given time.

Workaround: None.

VMware Cloud Director Container Service Extension fails to deploy clusters with TKG templates that have an unmodifiable placement policy set on them

Workaround: Perform the following tasks.

1. Log in to the VMware Cloud Director Tenant Portal as an administrator.
2. Click **Libraries** > **vApp Templates**.
3. In the **vApp Templates** window, select the radio button to the left of the template.
4. In the top ribbon, click **Tag with Compute Policies**.
5. Select the **Modifiable** check boxes, and click **Tag**.

In VMware Cloud Director 10.4, service providers cannot log in to the virtual machine of VMware Cloud Director Container Service Extension

In VMware Cloud Director 10.4, after deploying the VMware Cloud Director Container Service Extension virtual machine from OVA file, the following two check boxes in the VM settings page are not selected by default.

- Allow local administrator password
- Auto-generate password

Workaround: To allow service providers to log-in to the virtual machine of VMware Cloud Director Container Service Extension and perform troubleshooting tasks, select the Allow local administrator password and Auto-generate password check boxes.

1. Log in to VMware Cloud Director UI as a service provider and create a vApp from the VMware Cloud Director Container Service Extension OVA file.
2. After you deploy the vApp, and before you power it on, browse to **VM details** > **Guest OS Customization** and select **Allow local administrator password** and **Auto-generate password**.
3. After the update task finishes, power on the vApp.

Fast provisioned disks in Organization VCD cannot be resized

Workaround: To resize disks, deactivate fast provisioning in Organization VDC.

1. Log in to VMware Cloud Director UI as a provider, and select **Resources.**

2. In the **Cloud Resources** tab, select **Organization VDCs**, and select an organization VDC.
3. In the organization VDC window, under **Policies**, select **Storage**.
4. Click **Edit**, and deactivate the **Fast provisioning** toggle.
5. Click **Save**.

After you log in as a service provider and upload the latest Kubernetes Container Clusters UI plug-in, the CSE Management tab is not displayed

If there are multiple activated Kubernetes Container Clusters UI plug-ins with the same name or id but different versions, the lowest version of the plug-in is used. Only the highest version of the Kubernetes Container Clusters UI plug-in must be active. For more information on managing plug-ins, see Managing Plug-Ins.

Workaround: Deactivate the previous Kubernetes Container Clusters UI plug-ins.

1. Log in to **VMware Cloud Director** UI as a provider, and select **More > Customize Portal.**
2. Select the check box next to the names of the target plug-ins, and click **Enable** or **Disable**.
3. To start using the newly activated plug-in, refresh the Internet browser page.

Resize or upgrade a Tanzu Kubernetes Grid cluster by using kubectl

After a cluster is created in the Kubernetes Container Clusters UI plug-in, you can resize, upgrade, lifecycle manage the cluster, or manage workloads, by using `kubectl` instead of the Kubernetes Container Clusters UI plug-in.

1. To delete the RDE-Projector operator from the cluster, run `kubectl delete deployment -n rdeprojector-system rdeprojector-controller-manager`.
2. Detach the Tanzu Kubernetes Grid cluster from Kubernetes Container Clusters UI plug-in.
   a. In the VMware Cloud Director UI, in the **Cluster Overview** page, retrieve the cluster ID of the cluster.
   b. Update the RDE and set the `entity.spec.vcdKe.isVCDKECluster` value to false.
      a. To get the payload of the cluster, run `GET https://<vcd>/cloudapi/1.0.0/entities/<Cluster ID>`.
      b. Copy and update the json path in the payload.
      c. Set the `entity.spec.vcdKe.isVCDKECluster` value to false.
      d. Run `PUT https://<vcd>/cloudapi/1.0.0/entities/<Cluster ID>` with the modified payload. It is necessary to include the entire payload as the body of PUT operation.

      NOTE

      After performing the tasks above, the cluster is detached from VMware Cloud Director Container Service Extension 4.1 and you cannot manage the cluster through VMware Cloud Director Container Service Extension 4.1. You must use kubectl to manage, resize or upgrade the cluster by directly applying the cluster API specification, CAPI yaml.

VMware Cloud Director Container Service Extension 4.2 Release Notes

This document contains the following sections

Introduction

VMware Cloud Director Container Service Extension 4.2 | January 18 2024 | Build: 23133984
Check for additions and updates to these release notes.

What's New in January 2024

- A new version v0.1.2 of the `getting-started-airgapped` container is available to configure the private registry for airgapped functionality in VMware Cloud Director Container Service Extension 4.2. For more information, see Set Up a Local Container Registry in an Airgapped Environment.
- A new version v0.1.2 of the `cluster-upgrade-script-airgapped` container is available. You can use this to update the Kubernetes components in the pre-existing cluster to recommended versions. For more information, see Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters.
- Compatibility with VMware Cloud Director Extension for VMware Tanzu Mission Control 1.0. For more information, see VMware Cloud Director Extension for VMware Tanzu Mission Control.
- **Additional Tanzu Kubernetes Grid and Kubernetes support**: It is now possible to use the following Tanzu Kubernetes Grid and Kubernetes versions with VMware Cloud Director Container Service Extension 4.2:
  - Tanzu Kubernetes Grid 2.3.1 with Kubernetes 1.24, 1.25, and 1.26.
  - Tanzu Kubernetes Grid 2.4 with Kubernetes 1.25, 1.26, and 1.27.
- The **Auto Repair on Errors** toggle is deprecated, and will be unsupported starting with the next major VMware Cloud Director Container Service Extension release.
- The VMware Cloud Director IP spaces feature is available to use with VMware Cloud Director Container Service Extension 4.2. It is the responsibility of the service provider to add the rights to use this new feature to **Kubernetes Cluster Author** role and **Kubernetes Clusters Rights Bundle**. For more information see Kubernetes Cluster Author Role and Kubernetes Clusters Rights Bundle**.** To avail of this feature successfully, it is necessary to use Kubernetes Cloud Provider for VMware Cloud Director 1.6, and Kubernetes Cluster API Provider for VMware Cloud Director v1.3.0.
- Online and Offline Volume Expansion is now supported on disks provisioned by Kubernetes Container Storage Interface Driver for VMware Cloud Director 1.6.0. The feature needs to be activated manually by the Kubernetes Cluster Author. To activate this feature, the Kubernetes Cluster Author must update the `storageClass` with the `allowVolumeExpansion` flag activated as described here.

Upgrade

Kubernetes Container Clusters UI Plug-in 4.2 for VMware Cloud Director

A new version of Kubernetes Container Clusters UI plug-in is now available to use with VMware Cloud Director.

It is necessary to upgrade the Kubernetes Container Clusters UI plug-in before you upgrade the VMware Cloud Director Container Service Extension server.

The following steps outline how to upgrade the Kubernetes Container Clusters UI plug-in from 4.1.0 or 4.1.1 to 4.2:

1. Download the Kubernetes Container Clusters UI plug-in 4.2.
2. In the VMware Cloud Director Portal, from the top navigation bar, select **More > Customize Portal**.
3. Select the check box next to Kubernetes Container Clusters UI plug-in 4.1.0 or 4.1.1, and click **Disable**.
4. Click **Upload > Select plugin file**, and upload the Kubernetes Container Clusters UI plug-in 4.2 file.
5. Refresh the browser to start using the new plug-in.

For more information, refer to Managing Plug-Ins.

VMware Cloud Director Container Service Extension Server 4.2.

Service providers can now upgrade the VMware Cloud Director Container Service Extension Server to 4.2 through **CSE Management > Server Details > Update Server** in Kubernetes Container Clusters UI plug-in of VMware Cloud Director using the **Minor Version Upgrade** workflow.

For instructions on how to upgrade the VMware Cloud Director Container Service Extension Server from 4.1.0 or 4.1.1a to 4.2, see Minor Version Upgrade.

Compatibility Updates

VMware Cloud Director Container Service Extension 4.2 Interoperability Updates with Kubernetes Resources

To view the interoperability of VMware Cloud Director Container Service Extension 4.2 and previous versions with VMware Cloud Director, and additional product interoperability, refer to the Product Interoperability Matrix.

The following table displays the interoperability between VMware Cloud Director Container Service Extension 4.2, and Kubernetes resources.

Kubernetes Resources	Supported Versions	Documentation
Kubernetes Cloud Provider for VMware Cloud Director™	1.5.0	Kubernetes Cloud Provider for VMware Cloud Director Documentation
Kubernetes Container Storage Interface Driver for VMware Cloud Director™	1.5.0	https://github.com/vmware/cloud-director-named-disk-csi-driver#container-storage-interface-csi-driver-for-vmware-cloud-director-named-independent-disks
Kubernetes Cluster API Provider for VMware Cloud Director™	v1.2.0	https://github.com/vmware/cluster-api-provider-cloud-director
RDE Projector	0.7.0	Not applicable

Service providers can manually update Kubernetes resources through the following workflow:

1. In VMware Cloud Director UI, from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. In Kubernetes Container Clusters UI plug-in 4.2, select **CSE Management** > **Server Details** > **Update Server** > **Update Configuration** > **Next**.
3. In the **Current CSE Server Components** section, update the Kubernetes resources configuration.
4. Click **Submit Changes**.

For more information, see Update the VMware Cloud Director Container Service Extension Server.

In VMware Cloud Director Container Service Extension, it is necessary to confirm that the Kubernetes components in a cluster have the required versions.

For clusters that were created using older versions of VMware Cloud Director Container Service Extension, it is necessary to perform a one time script upgrade action. This allows the clusters to be compatible with the latest VMware Cloud Director Container Service Extension.

The Kubernetes Container Clusters UI warns users about this requirement with the following warning when tenant users attempt certain workflows in the UI:

```
Confirm that the components in this cluster have the required versions.
```

You can ignore this UI warning if the relevant Kubernetes component versions in the **Cluster Information** page in Kubernetes Container Clusters UI match the supported versions displayed in the above table. If this warning appears, and the current versions of Kubernetes components in the cluster do not match the available versions, follow the instructions in the Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters. Do not continue with the workflow you are currently in.

VMware Cloud Director Container Service Extension 4.2 does not support Tanzu Kubernetes Grid versions 1.6.1, 1.5.4 and 1.4.3.

Tanzu Kubernetes Grid versions 1.6.1, 1.5.4 and 1.4.3 are no longer supported by VMware. For more information on the end of this support, see the lifecycle matrix.

It is necessary for service providers and tenant users to upgrade pre-existing Tanzu Kubernetes Grid 1.6.1, 1.5.4 and 1.4.3 clusters to Tanzu Kubernetes Grid versions 2.1.1, 2.2, 2.3.1 or 2.4 and supported Kubernetes versions. New cluster deployment attempts from VMware Cloud Director Container Service Extension 4.2 using unsupported Tanzu Kubernetes Grid versions 1.6.1, 1.5.4 and 1.4.3 will fail.

Resolved Issues

For API users, the cluster list page displays an error message if the node pool names in the cluster's capi yaml are changed after cluster creation.

The cluster list datagrid fails to load, and displays the following error message:

```
Error: Failed to fetch Kubernetes clusters
```

This issue is fixed for VMware Cloud Director Container Service Extension 4.2.

For API users, the cluster list page displays an error message if the control plane node pool name does not end in `-control-plane-node-pool.`

The cluster list datagrid fails to load, and displays the following error message:

```
Error: Failed to fetch Kubernetes clusters
```

This error means that API users must name their control plane node pool in a format that ends with `-control-plane-node-pool.`

This issue is fixed for VMware Cloud Director Container Service Extension 4.2.

In the Kubernetes Container Clusters UI plugin, new worker node pools are created using parameters that were not specified in the Create New Worker Node Pools menu.

Unspecified values in the **Create New Worker Node Pools** menu may result in the node pool using the parameters of other node pools. This issue affects the following parameters:

- vGPU Activated
- Sizing Policy
- Placement Policy or vGPU Policy (if GPU toggle is activated)
- Storage Profile
- Disk Size

This issue is fixed for VMware Cloud Director Container Service Extension 4.2.

In the Kubernetes Container Clusters UI plugin, the Kubernetes Version dropdown menu in the cluster creation wizard for a TKGs cluster displays a spinner indefinitely.

On the **Kubernetes Policy** page in the cluster creation wizard, the **Kubernetes Version** dropdown selection displays a spinner indefinitely.

This occurs due to the supported Kubernetes version API sending an invalid response to the Kubernetes Container Clusters UI plugin, so the UI plugin fails to parse it.

This issue is fixed for VMware Cloud Director Container Service Extension 4.2.

GPU operator installation fails on the vGPU node in the cluster.

The Kubernetes Cluster API Provider for VMware Cloud Director component previously installed `nvidia-container-runtime` package, and edited `containerd` runtime to point to `nvidia-container-runtime` during the provisioning of vGPU node in the VMware Cloud Director Container Service Extension Kubernetes cluster. This customization interfered with Nvidia GPU Operator installation.

This issue is fixed in Kubernetes Cluster API Provider for VMware Cloud Director v1.1.1.

Kubernetes Cluster API Provider for VMware Cloud Director v1.1.1 removes the customization in favor of GPU Operator installation. Users who want to manually install drivers on vGPU node are required to perform the following customization workflow:

1. Install the `nvidia-container-runtime` package
2. Update the `containerd` configuration to point to the NVIDIA runtime.

For more information. see Kubernetes Cluster API Provider for VMware Cloud Director v1.1.1 Release Notes.

The CSE Management workflow in a multi-site VMware Cloud Director setup only allows for a single server configuration entity.

In Kubernetes Container Clusters UI plug-in, the **CSE Management** workflow in a multi-site VMware Cloud Director setup may display a server configuration entity that belongs to a different site. This results in the **CSE Management** workflows failing in multi-site environments.

This issue is fixed in VMware Cloud Director Container Service Extension 4.1.1a. The **CSE Management** workflow in Kubernetes Container Clusters UI plug-in now only fetches the server configuration entity that belongs to the site where the user is currently logged-in. This fix allows each site in a multi-site environment to create, and manage its own server configuration entity.

In VMware Cloud Director Service Provider Portal, if a service provider navigates into a specific cluster, and returns to the landing page, the cluster list does not display all clusters.

Service providers can view a full list of clusters that are in an environment in VMware Cloud Director Service Provider portal. From this view, if a service provider clicks in to a cluster to view details of that cluster, navigates to the **Persistent Volumes** tab, and clicks back to return to the listing of all clusters, then the original list of all the clusters is not visible. Some clusters do not display on the list.

This issue is fixed for VMware Cloud Director Container Service Extension 4.1.1a.

The Kubernetes Container Clusters UI plug-in does not display the Container Registry setting in Server Details page.

In Kubernetes Container Clusters UI plug-in, in the **CSE Management** tab, the **Server Details** tab does not display the **Container Registry** setting.

This issue is fixed in VMware Cloud Director Container Service Extension 4.1.1a.

The Projector version is not updated in RDE after Projector deployment is upgraded in the cluster.

The version of the Projector component is not updated in the cluster RDE's projector status section, even though the version in the Projector deployment changes.

This issue is fixed in VMware Cloud Director Container Service Extension 4.1.1a.

vApp creation failure events does not display error message

During cluster creation, if the create vApp operation fails then the **Event Details** tab for that cluster in the Kubernetes Container Clusters UI-Plugin does not show full error message.

This issue is fixed in VMware Cloud Director Container Service Extension 4.1.1a. The error message is now included in the **Detailed Error** section that aids in troubleshooting.

Tanzu Standard Repository is not installed with clusters that are created in VMware Cloud Director Container Service Extension 4.1 with Tanzu Kubernetes Grid 2.1.1 and 2.2.

In VMware Cloud Director Container Service Extension 4.1.1a, clusters created with Tanzu Kubernetes Grid 2.1.1 and 2.2 automatically have the Tanzu-standard repository installed.

For clusters that were created using VMware Cloud Director Container Service Extension 4.1, it is necessary to install the repository and packages.

Upgrading a cluster using Kubernetes Container Clusters UI plugin 4.1 fails when the cluster was initially created using Kubernetes Container Clusters UI plugin 4.0.

When you use Kubernetes Container Clusters UI plugin 4.1 to upgrade a cluster that was created in Kubernetes Container Clusters UI plugin 4.0, the following error is seen in the **Events** tab, even though the cluster-upgrade-script was executed successfully without any errors:

```
Error: PatchObjectError
error message: [VCDCluster.infrastructure.cluster.x-k8s.io "<cluster-name>" is invalid: spec.loadBalancerCon-
figSpec: Invalid value: "null": spec.loadBalancerConfigSpec in body must be of type object: "null"] during
 patching objects
```

This issue is fixed in VMware Cloud Director Container Service Extension 4.1.1a.

Auto Repair on Errors toggle must be deactivated immediately after a cluster is created.

It is possible that a provisioned cluster can go into an error state due to a known issue. If the **Auto Repair on Errors** feature is activated on the cluster, that cluster can get deleted and recreated, which causes disruption of workloads on that cluster.

When you create clusters, it is recommended to deactivate the **Auto Repair on Errors** toggle to avoid clusters from getting deleted, and recreated if they go into error state.

For more information on the **Auto Repair on Errors** in the cluster creation workflow, see Create a VMware Tanzu Kubernetes Grid Cluster.

> **NOTE**
>
> The **Auto Repair on Errors** setting is deactivated by default in the Kubernetes Container Clusters UI. If you activate it for any reason, you must turn it off immediately after cluster is provisioned.

This issue has been fixed for VMware Cloud Director Container Service Extension 4.1.1a.

VMware Cloud Director Container Service Extension Server additionally deactivates the **Auto Repair on Error** toggle after a cluster is created successfully.

When new organizations are added to VMware Cloud Director, the VMware Cloud Director Container Service Extension server may fail to provide access to VMware Cloud Director Container Service Extension configuration to the new organizations.

When this issue occurs, the following error message appears in the VMware Cloud Director Container Service Extension log file:

```
"msg":"error occurred while onboarding new tenants with ReadOnly ACLs for VCDKEConfig:
[unable to get all orgs: [error occurred retrieving list of organizations: [error getting
list of organizations: 401 Unauthorized]]]"
```

Also, tenant users may see the following warning message and be blocked when they try to create a cluster using VMware Cloud Director Container Service Extension UI.

```
Cannot fetch provider configuration. Please contact your administrator
```

This issue is fixed for VMware Cloud Director Container Service Extension 4.1.1a.

For each cluster, repeated messages of `Invoked getFullEntity (urn:vcloud:entity:vmware:capvcdCluster:{ID})` display in the Recent Tasks pane of VMware Cloud Director UI.

This issue is happening because VMware Cloud Director Container Service Extension is retrieving RDE for all the clusters instead of retrieving RDE for unprocessed clusters.

This issue has been fixed for VMware Cloud Director Container Service Extension 4.1.1a.

Clusters created using Kubernetes Cluster API Provider for VMware Cloud Director (CAPVCD) management cluster, without involvement of VMware Cloud Director Container Service Extension server, displays a Pending status in Kubernetes Container Clusters UI.

This issue has been fixed for VMware Cloud Director Container Service Extension 4.1. The status of these clusters is now **Non-CSE**. The only permitted operation for these non-VMware Cloud Director Container Service Extension clusters is to download the `kube config` of the cluster.

The Kubernetes Container Clusters UI plugin storage profile selection form fields do not filter storage policies by `entitytype`.

The storage profile selection form fields display all storage profiles visible to the logged-in user, such as VMs, vApps, Catalog items, or named disks.

This issue has been fixed for VMware Cloud Director Container Service Extension 4.1. In Kubernetes Container Clusters UI in VMware Cloud Director Container Service Extension 4.1, the storage policy selection only shows storage policies that support any of these entitytypes:

- vApp and VM templates
- Virtual machines

Kubernetes cluster resize operation fails in VMware Cloud Director Container Service Extension 4.0.x.

If users attempt to change organization VDC names in VMware Cloud Director after clusters are created, further cluster operations such as cluster resize can fail.

This issue has been fixed for VMware Cloud Director Container Service Extension 4.1.

When a node of the cluster is deleted due to failure in vSphere or other underlying infrastructure, VMware Cloud Director Container Service Extension does not inform the user, and it does not auto-heal the cluster.

When the node of a cluster is deleted, basic cluster operations, such as cluster resize and cluster upgrade, continue to work. The deleted node remains in deleted state, and is included in computations regarding size of the cluster.

This issue has been fixed for VMware Cloud Director Container Service Extension 4.1, as the Node Health Check Configuration feature addresses this occurence.

The cluster creation for multi-control plane or multi-worker node goes into an error state. The Events tab in the cluster details page shows an `EphemeralVMError` event due to the failure to delete ephemeralVM in VMware Cloud Director.

The same error events can appear repeatedly if the **Auto Repair on Errors** setting is activated on the cluster. If the **Auto Repair on Errors** setting is off, sometimes the cluster can show an error state due to the failure to delete the ephemeralVM in VMware Cloud Director even though the control plane and worker nodes are created successfully.

This issue is visible in any release and patch release after but not including VMware Cloud Director 10.3.3.3, and any release and patch release starting with VMware Cloud Director 10.4.1.

This issue is fixed for VMware Cloud Director Container Service Extension 4.1 release.

When a force delete attempt of a cluster fails, the `ForceDeleteError` that displays in the Events tab of the cluster info page does not provide sufficient information regarding the failure to delete the cluster.

This issue is fixed for VMware Cloud Director Container Service Extension 4.1 release.

Known Issues

In VMware Cloud Director Container Service Extension 4.2, if you force delete clusters that are configured to use IP Spaces, the IP allocated to the cluster and/or Kubernetes services running on the cluster are not released automatically, and manual intervention is necessary.

Creating clusters with Node Health Check activated, or activating Node Health Check for existing clusters causes the cluster to become unmanageable.

This error is caused by Kubernetes Container Clusters UI plugin 4.2 creating the **MachineHealthCheck** capi yaml object with an invalid **apiVersion** value. Clusters that contain this invalid **MachineHealthCheck** section become unmanageable.

There are two ways in Kubernetes Container Clusters UI 4.2 plugin for a cluster to enter this unmanageable state:

- If a user creates a cluster with **Node Health Check** activated, then the resulting cluster will be unmanageable. Users should ensure that **Node Health Check** is deactivated when creating clusters with UI plugin 4.2.
- If a user activates **Node Health Check** for a cluster that has never previously activated **Node Health Check**, the cluster will become unmanageable.

  NOTE

  If a cluster was created using a lower UI plugin version and **Node Health Check** had already been activated before, then that cluster can activate or deactivate **Node Health Check** without encountering this issue.

Workaround:

The only workaround for the Kubernetes Container Clusters UI Plugin 4.2 to fix a cluster that has this issue is to manually update the cluster's **MachineHealthCheck apiVersion** value to **cluster.x-k8s.io/v1beta1** through API.

The Create a Tanzu Kubernetes Grid Cluster workflow, and Create New Worker Node Pools workflow fails when sizing policy, and vGPU policy are both specified if the vGPU policy already contains sizing information.

Workaround:

If you select a vGPU policy that already contains sizing information during **Create a Tanzu Kubernetes Grid Cluster** workflow, or **Create New Worker Node Pools** workflow, do not also select a sizing policy.

The Upgrade Availability value on the cluster list page, and on the Cluster Information page displays `Error` .

When you select a cluster in the cluster list datagrid, `Error` displays in the **Upgrade** column.

In the **Cluster Information** page, the **Upgrade Availability** property displays `Error` .

Workaround:

This occurs if there is a corrupted or invalid vApp template in a catalog that is visible to the user. Remove the problematic vApp template from the catalog, and then the user must refresh their browser.

Kubernetes templates datagrid displays error message, and fails to load.

When the Kubernetes templates datagrid fails to load, the following error message displays:

```
Error: Failed to fetch Kubernetes Templates
```

Workaround:

This occurs if there is a corrupted or invalid vApp template in a catalog that is visible to the user. Remove the problematic vApp template from the catalog, and then the user must refresh their browser.

`java.lang.OutOfMemoryError: Java heap space` error causes VMware Cloud Director services to fail continuously after startup.

VMware Cloud Director crashes due to `OutOfMemoryError` . The issue occurs when the resolve operation is invoked on an RDE that has a lot of tasks associated with it. The issue is present on VMware Cloud Director 10.4 and above. For more information, see VMware Knowledge Base Article 95464.

Do not change the Registry URL in VMware Cloud Director Container Service Extension configuration as changes are not supported.

Use load balancers to front registry virtual machines to swap the virtual machines out if necessary.

If you use VMware Cloud Director 10.4.2.2, the cluster deletion workflow can fail in Kubernetes Container Clusters UI.

The cluster deletion operation fails with the following error:

```
"error": "failed to delete VCD Resource [clusterName] of type [VApp] from VCDResourceSet
of RDE [urn:vcloud:entity:vmware:capvcdCluster:<uuid>]: [failed to update capvcd status
for RDE [urn:vcloud:entity:vmware:capvcdCluster:<uuid>]; expected http response [200],
obtained [400]: resp: [\"{\\\"minorErrorCode\\\":\\\"BAD_REQUEST\\\",\\\"message\\\":\\
\"[ a8e89bd2-195d-458b-808d-3ff81e074fa0 ] RDE_CANNOT_VALIDATE_AGAINST_SCHEMA [ #/status/
capvcd/vcdResourceSet/2: expected type: JSONObject, found: Null\\\\n ]\\\",\\\"stackTrace\
\\":null}\"]: [400 Bad Request]]"
```

This is a bug of VMware Cloud Director 10.4.2.2.

Workaround: Use **Force Delete** workflow to delete the cluster.

In Kubernetes Container Clusters UI plug-in, the CSE Management upgrade workflows may add or remove rights from the CSE Admin Role or Kubernetes Cluster Author.

It is necessary to manually update the Custom roles that are cloned from CSE Admin Role or Kubernetes Cluster Author. If you do not do this, users can face errors during cluster workflows.

VMware Cloud Director Container Service Extension does not automaticially install a Tanzu-standard repository in the Tanzu Kubernetes Grid 2.1.1 and 2.2 clusters.

Workaround:

Install the repository and packages using documentation for Tanzu Kubernetes Grid 2.2 and 2.1.1 clusters.

Tenant users may see the following warnings when they attempt certain actions using VMware Cloud Director Container Service Extension.

Warnings:

- ***Cannot fetch provider configuration. Please contact your administrator.***
  Tenant users may see this warning, and be blocked when they try to create a cluster.
- ***Node Health Check settings have not been configured by your provider.***
  Tenant users may see this warning when they try to activate **Node Health Check** during cluster creation or in the cluster settings.

These warnings can occur for the following reasons:

- The VMware Cloud Director Container Service Extension server has not finished starting up.
- The VMware Cloud Director Container Service Extension server has not yet published the server configuration to tenant organizations. The server configuration is published automatically every hour from the server startup as the server is running. Therefore, publishing to new tenant organizations that are created during hourly window occurs at the end of the hour.
- The tenant user's role does not have the following right: **View: VMWARE:VCDKECONFIG**. This right was added to the Kubernetes Cluster Author global role in VMware Cloud Director Container Service Extension 4.1.
- There was an unexpected error while fetching the server configuration.

Workaround:

- Service providers must ensure the VMware Cloud Director Container Service Extension server is set up, and operating successfully.
- Ensure the tenant user's role has the right **View: VMWARE:VCDKECONFIG.** Tenant users must log out of VMware Cloud Director, and log back in to activate any changes made to their role.
- Wait for hourly publishing to new organizations.

In some instances, nodes cannot join clusters even when the cluster is in an available state. This issue can occur intermittently.

The following error appears in the **Events** tab of the cluster info page in Kubernetes Container Clusters UI:

`VcdMachineScriptExecutionError` with the following details:

```
script failed with status [x] and reason [Date Time 1 /root/node.sh: exit [x]]
```

Workaround:

For VMware Cloud Director Container Service Extension 4.1, there is a retry mechanism added that uses a retry feature from Cluster API which reduces the occurrence of this issue.

VMware Cloud Director Container Service Extension 4.1 does not support Dynamic Host Configuration Protocol (DHCP)

The cluster creation workflow in VMware Cloud Director Container Service Extension 4.1 fails if the cluster is connected to a routed organization VDC network that uses DHCP instead of static IP pool to distribute IPs to virtual machines.

VMware Cloud Director Container Service Extension 4.1 only supports organization VDC networks in the following circumstances:

- If the VDC is routed.
- If the VDC uses static IP pool to distribute IPs to virtual machines that are connected to it.

It is not possible to activate GPU support in an airgapped cluster.

As VMware cannot redistribute nVidia packages, it is not possible to activate GPU support in an airgapped cluster out of box. The failure occurs when the cluster attempts to download the nVidia binary from `nvidia.github.io` in the cloud initialization script.

Workaround:

As VMware cannot redistribute nVidia packages, it is not possible to activate GPU support in an airgapped cluster out of box. The failure occurs when the cluster attempts to download the nVidia binary from `nvidia.github.io` in the cloud initialization script. Service providers can potentially consider allowing the cluster access to `nvidia.github.io` by using a proxy server.

`Audit_trail` table grows rapidly in the VMware Cloud Director database due to RDE modify events being too large.

RDE modify events log the whole body of the RDE that has changed. These large events cause the `audit_trail` table to grow longer than necessary.

Workaround:

Upgrade to VMware Cloud Director 10.3.3.4 or above. If you are using VMware Cloud Director 10.3.3.4, set the `audit.rde.diffOnly` config property to `True` .

If you are using VMware Cloud Director 10.4.0 and above, there is no requirement to set any configuration property.

VMware Cloud Director Container Service Extension 4.1 uses Kubernetes Cluster API Provider for VMware Cloud Director 1.1 and Kubernetes Cloud Provider for VMware Cloud Director 1.4 as default. These two component versions do not support IP spaces.

Tanzu Addons-Manager does not appear after upgrading to Tanzu Kubernetes Grid 2.2.0 with Kubernetes v1.24+.

After you upgrade a VMware Cloud Director Container Service Extension 4.0.3 cluster from Tanzu Kubernetes Grid 1.6.1 with Kubernetes v1.23.x to Tanzu Kubernetes Grid 2.2 with Kubernetes v1.24.x, `tanzu-addon-controller-manager` pod is stuck at **PENDING** or `CrashLoopBackOff` state for the following reason:

```
Error from server (NotFound): packageinstalls.packaging.carvel.dev "addons-
manager.tanzu.vmware.com" not found
```

Use the following workaround to manually delete the `tanzu-addons-controller-manager` deployment and PackageInstall object.

To delete the deployment, perform the following commands:

```
kubectl get deployments -A
kubectl delete deployment -n tkg-system   tanzu-addons-controller-manager
```

To delete the `PackageInstall` object, perform the following commands:

```
kubectl get packageinstall -A
kubectl delete packageinstall -n tkg-system   tanzu-addons-manager
```

An ephemeral VM is created during the cluster creation process, and is deleted by VMware Cloud Director Container Service Extension when the cluster creation process is complete. It is possible that the API request to delete the ephemeral VM can fail.

VMware Cloud Director Container Service Extension reattempts to delete the ephemeral VM for up to 15 minutes. In an event that VMware Cloud Director Container Service Extension fails to delete the ephemeral VM after reattempting, it leaves the ephemeral VM in the cluster's VApp without deleting it.

The following error appears in the **Events** tab of the cluster info page in Kubernetes Container Clusters UI:

`EphemeralVMError` with the following details:

```
error deleting Ephemeral VM [EPHEMERAL-TEMP-VM] in vApp [cluster-vapp-name]: [reason for
failure]. The Epemeral VM needs to be cleaned up manually.
```

The reason for failure depends on the stage at which the ephemeral VM deletion failed. Once you observe this notification, it is safe to delete the ephemeral VM from the cluster's VApp in the VMware Cloud Director UI.

Workaround:

1. Log in to the VMware Cloud Director Tenant Portal, and from VMware Cloud Director navigation menu, select **Data Centers**.
2. In the Virtual Data Center page, select the organization tile, and from the left navigation menu, select **vApps.**
3. In the vApps page, select the vApp of the cluster.
4. In the cluster information page, click the ellipse to the left of the Ephermal VM, and click **Delete**.

However, if the ephemeral VM is not manually cleaned up, and if a delete request is issued, the cluster delete operation fails. It is then necessary to force delete the cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster, and in the cluster information page, click **Delete**.
3. In the **Delete Cluster** page, select the **Force Delete** checkbox, and click **Delete**.

It is not possible to create clusters in VMware Cloud Director Container Service Extension 4.1. when using a direct organization VDC network with NSX in VMware Cloud Director.

VMware Cloud Director Container Service Extension 4.1 clusters do not support this configuration.

In VMware Cloud Director Container Service Extension, the creation of Tanzu Kubernetes Grid clusters can fail due to a script execution error.

The following error appears in the **Events** tab of the cluster info page in Kubernetes Container Clusters UI:

`ScriptExecutionTimeout` with the following details:

```
error while bootstrapping the machine [cluster-name/EPHEMERAL_TEMP_VM]; timeout for post
customization phase [phase name of script execution]
```

Workaround:

When this error occurs, it is recommended to activate **Auto Repair on Errors** from cluster settings. This instructs VMware Cloud Director Container Service Extension to reattempt cluster creation.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster, and in the cluster information page, click **Settings,** and activate the **Auto Repair on Errors** toggle.
3. Click **Save.**

   NOTE

   It is recommended to deactivate the **Auto Repair on Errors** toggle when troubleshooting cluster creation issues.

In Kubernetes Container Clusters UI plug-in, the cluster delete operation can fail when the cluster status is Error.

To delete a cluster that is in **Error** status, it is necessary to force delete the cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters.**
2. Select a cluster, and in the cluster information page, click **Delete**.
3. In the **Delete Cluster** page, select the **Force Delete** checkbox, and click **Delete**.

`ERROR: failed to create cluster: failed to pull image` failure

This error occurs in the following circumstances:

- When a user attempts to create a Tanzu Kubernetes Grid Cluster using VMware Cloud Director Container Service Extension 4.1, and it fails intermittently.
- An image pull error due to a HTTP 408 response is reported.

This issue can occur if there is difficulty reaching the Internet from the EPHEMERAL_TEMP_VM to pull the required images.

Potential causes:

- Slow or intermittent Internet connectivity.
- The network IP Pool cannot resolve DNS (docker pull error).
- The network MTU behind a firewall must set lower.

To resolve the issue, ensure that there are no networking connectivity issues stopping the EPHEMERAL_TEMP_VM from reaching the Internet.

For more information, refer to https://kb.vmware.com/s/article/90326.

Users may encounter authorization errors when executing cluster operations in Kubernetes Container Clusters UI plug-in if a Legacy Rights Bundle exists for their organization.

- After you upgrade VMware Cloud Director from version 9.1 or earlier, the system may create a **Legacy Rights Bundle** for each organization. This **Legacy Rights Bundle** includes the rights that are available in the associated organization at the time of the upgrade and is published only to this organization. To begin using the rights bundles model for an existing organization, you must delete the corresponding **Legacy Rights Bundle**.
- In the **Administration** tab in the service provider portal, you can delete **Legacy Rights Bundles**. Kubernetes Container Clusters UI plug-in **CSE Management** has a server setup process that automatically creates, and publishes **Kubernetes Clusters Rights Bundle** to all tenants. The rights bundle contains all rights that are involved in Kubernetes cluster management in VMware Cloud Director Container Service Extension 4.0.

Resizing or upgrading a Tanzu Kubernetes Grid cluster using kubectl.

After a cluster has been created in the Kubernetes Container Clusters UI plug-in, you can use kubectl to manage workloads on Tanzu Kubernetes Grid clusters.

If you also want to lifecycle manage, resize and upgrade the cluster through kubectl instead of the Kubernetes Container Clusters UI plug-in, complete the following steps:

1. Delete the RDE-Projector operator from the cluster `kubectl delete deployment -n rdeprojector-system rdeprojector-controller-manager`
2. Detach the Tanzu Kubernetes Grid cluster from Kubernetes Container Clusters UI plug-in.
   a. In the VMware Cloud Director UI, in the **Cluster Overview** page, retrieve the cluster ID of the cluster.
   b. Update the RDE with `entity.spec.vcdKe.isVCDKECluster` to false.
      a. Get the payload of the cluster - GET `https://<vcd>/cloudapi/1.0.0/entities/<Cluster ID>`
      b. Copy and update the json path in the payload. – `entity.spec.vcdKe.isVCDKECluster` to false.
      c. PUT `https://<vcd>/cloudapi/1.0.0/entities/<Cluster ID>` with the modified payload. It is necessary to include the entire payload as the body of PUT operation.
   c. At this point the cluster is detached from VMware Cloud Director Container Service Extension 4.1, and it is not possible to manage the cluster through VMware Cloud Director Container Service Extension 4.1. It is now possible to use kubectl to manage, resize or upgrade the cluster by applying CAPI yaml, the cluster API specification, directly.

Policies selection in VMware Cloud Director Container Service Extension 4 plug-in does not populate the full list after selection for the purpose of policy modification.

When a user selects a sizing policy in the **Kubernetes Container Clusters** plug-in and they want to change it, the dropdown menu only displays the selected sizing policy, and does not automatically load alternative sizing policies.

The user has to delete the text manually to allow the alternative sizing policies to appear. This also occurs in the dropdown menu when the user selects of placement policies and storage policies.

This is intentional. This is how the combobox html, ***Clarity***, web component works.

Note: ***Clarity*** is the web framework that **VMware Cloud Director** UI is built on.

The dropdown box uses the input text as a filter. When nothing is in the input field, you can see all selections, and the selections filter as you type.

When you create a VMware Cloud Director Container Service Extension cluster, a character capitalization error appears.

In the **Kubernetes Container Clusters** UI, if you use capital letters, the following error appears:

- *Name must start with a letter, end with an alphanumeric, and only contain alphanumeric or hyphen (-) characters. (Max 63 characters)*

This is a restriction set by **Kubernetes**. Object names are validated under RFC 1035 labels. For more information, refer to Kubernetes website.

Kubernetes Container Clusters UI-Plugin 4.1 does not interoperate with other Kubernetes Container Clusters UI plug-ins, such as 4.0 and 3.5.0.

The ability to operate these two plug-ins simultaneously without conflict is a known VMware Cloud Director UI limitation. You can only have one plug-in activated at any given time.

VMware Cloud Director Container Service Extension fails to deploy clusters with TKG templates that have an unmodifiable placement policy set on them.

1. Log in to the VMware Cloud Director Tenant Portal as an administrator.
2. Click **Libraries** > **vApp Templates**.
3. In the **vApp Templates** window, select the radio button to the left of the template.
4. In the top ribbon, click **Tag with Compute Policies**.
5. Select the **Modifiable** checkboxes, and click **Tag**.

In VMware Cloud Director 10.4, service providers are unable to log-in to the VMware Cloud Director Container Service Extension virtual machine by default.

In **VMware Cloud Director** 10.4, after deploying the **VMware Cloud Director Container Service Extension** virtual machine from OVA file, the following two checkboxes in the VM settings page are not selected by default:

- Allow local administrator password
- Auto-generate password

It is necessary to select these checkboxes to allow providers to log-in to the **VMware Cloud Director Container Service Extension** virtual machine in future to perform troubleshooting tasks.

1. Log in to **VMware Cloud Director** UI as a service provider, and create a vApp from the **VMware Cloud Director Container Service Extension** OVA file.
2. Once you deploy the vApp, and before you power it on, go to **VM details** > **Guest OS Customization** > Select **Allow local administrator password** and **Auto-generate password**.
3. After the vApp update task finishes, power on the vApp.

Fast provisioning must be deactivated in Organization VDC in order to resize disks.

1. Log in to VMware Cloud Director UI as a provider, and select **Resources.**
2. In the **Cloud Resources** tab, select **Organization VDCs**, and select an organization VDC.
3. In the organization VDC window, under **Policies**, select **Storage**.
4. Click **Edit**, and deactivate the **Fast provisioning** toggle.
5. Click **Save**.

When you log in as a service provider, after you upload the latest UI plug-in, the CSE Management tab does not display.

Deactivate the previous UI plug-in that is built into **VMware Cloud Director**.

1. Log in to **VMware Cloud Director** UI as a provider, and select **More > Customize Portal.**
2. Select the check box next to the names of the target plug-ins, and click **Enable** or **Disable**.
3. To start using the newly activated plug-in, refresh the Internet browser page.

   NOTE

   If there are multiple activated plugins with the same name or id but different version, the lowest version plug-in is used. Therefore, only activate the highest version plug-in. Deactivate all other version plug-ins.

For more information on managing plug-ins, see Managing Plug-Ins.

Installing, Configuring, and Upgrading VMware Cloud Director Container Service Extension as a Service Provider

This guide provides information about how to install, configure, and upgrade VMware Cloud Director™ Container Service Extension™ software, and how to configure it to work with VMware Cloud Director™.

Intended Audience

This guide is intended for **Service Providers** who want to install, configure and upgrade VMware Cloud Director Container Service Extension software, and the VMware Cloud Director Container Service Extension server.

What is VMware Cloud Director Container Service Extension

VMware Cloud Director™ Container Service Extension™ is a VMware Cloud Director extension that offers a server and a UI component to tenant users to create Tanzu Kubernetes Grid clusters in their virtual data centers alongside their virtual machines.

VMware Cloud Director Container Service Extension brings Kubernetes as a service to VMware Cloud Director, offering multi-tenant, VMware supported, production ready, and compatible Kubernetes services with Tanzu Kubernetes Grid. As a service provider administrator, you can add the service to your existing VMware Cloud Director tenants. By using VMware Cloud Director Container Service Extension, customers can also use Tanzu products and services such as Tanzu® Mission Control to manage their clusters. For more information on Tanzu Mission Control, refer to Tanzu Mission Control.

VMware Cloud Director Container Service Extension brings Kubernetes as a service to VMware Cloud Director by deploying and managing fully functional VMware Cloud Director provisioned clusters. By using VMware Cloud Director Container Service Extension, development teams can focus on application development, and simplifies infrastructure management.

The VMware Cloud Director Container Service Extension server is highly available, maintaining availability continuously. This means that if an incident occurs, the VMware Cloud Director Container Service Extension server can shift workloads and configurations away from the affected nodes, and continue to operate. There is no requirement to backup or restore VMware Cloud Director Container Service Extension appliance as it is a stateless application. For VMware Cloud Director Container Service Extension server to achieve high availability, it is necessary for service providers to run multiple vApp instances of the server. For more information, see Create a vApp from VMware Cloud Director Container Service Extension server OVA file.

Figure 1: Architecture diagram of VMware Cloud Director Container Service Extension 4.x and workflows of service providers and tenant users

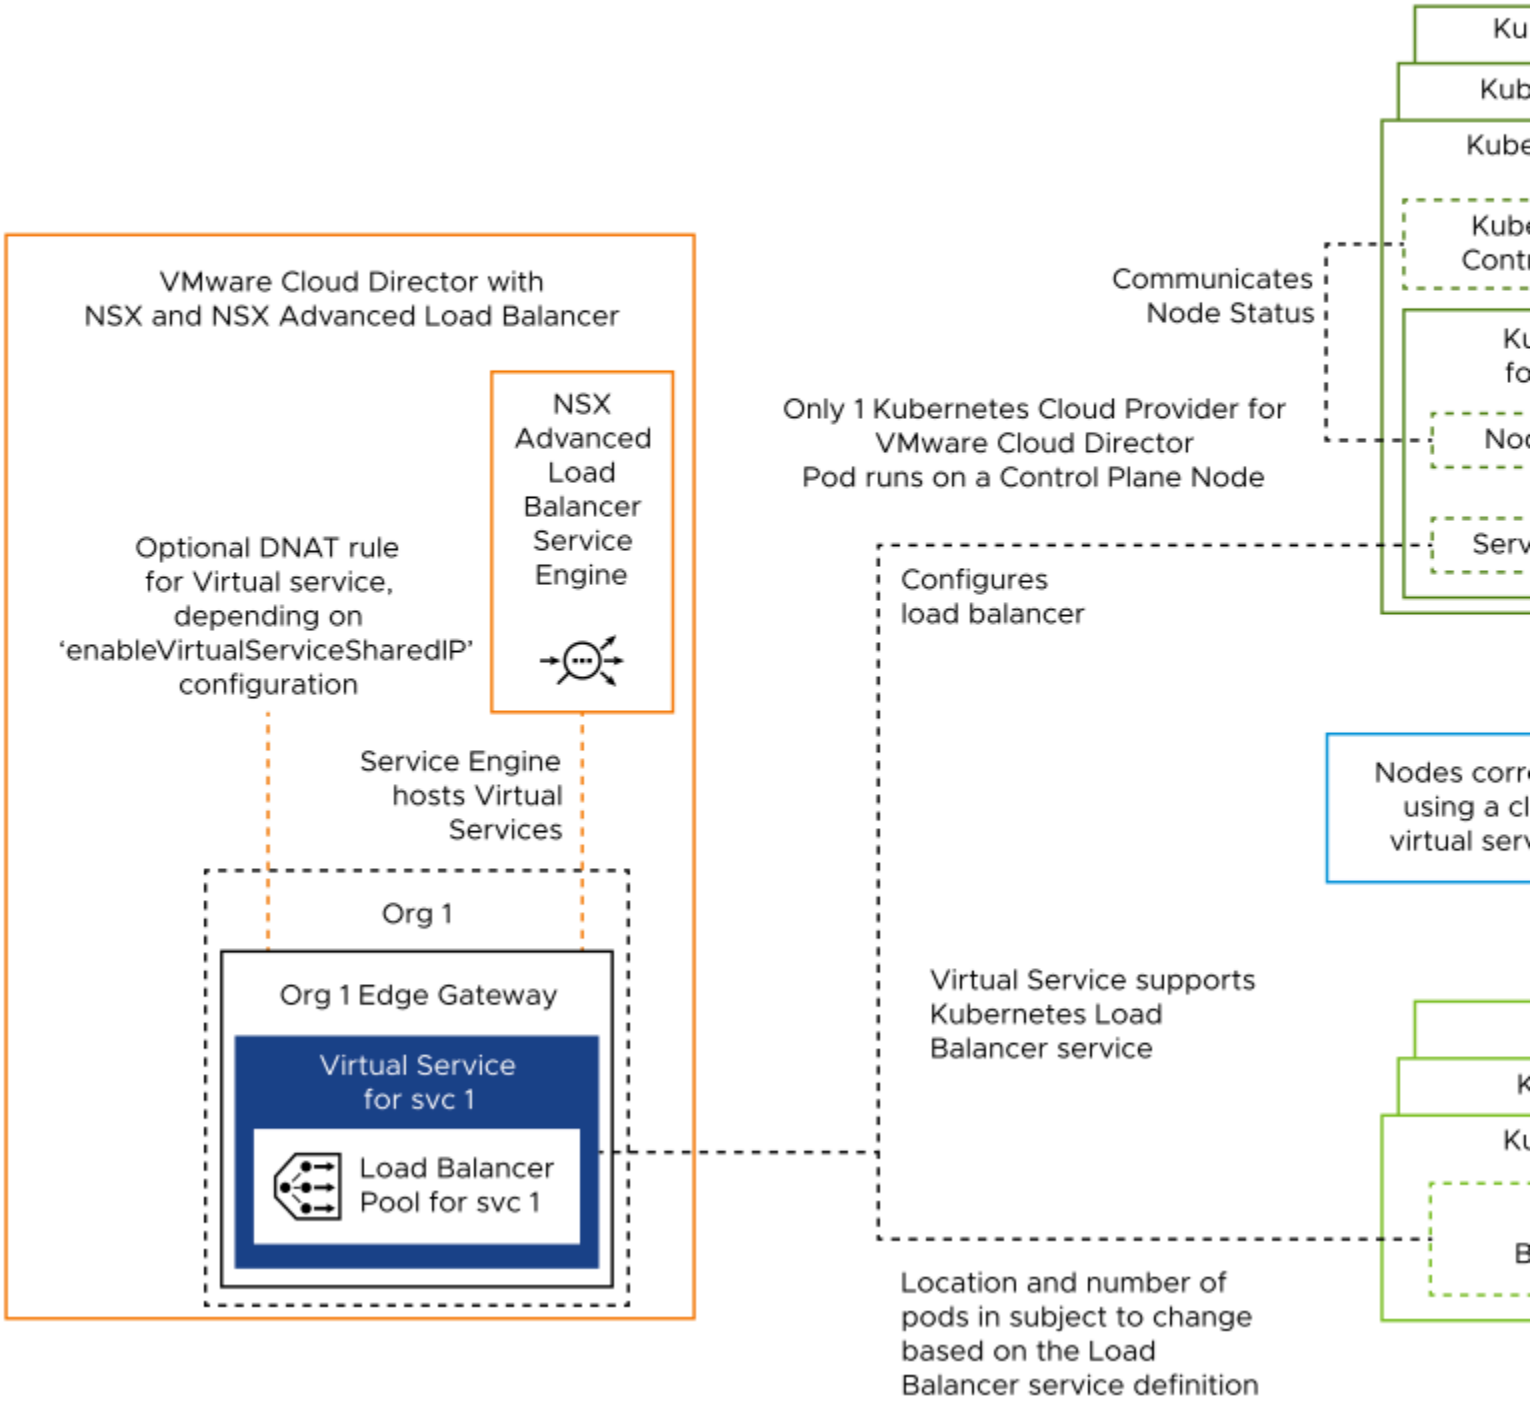

As a service provider administrator, you can perform the following functions in VMware Cloud Director Container Service Extension:

- Set up the VMware Cloud Director Container Service Extension server through the new **CSE Management** tab in the Kubernetes Container Clusters plug-in. You can also perform cluster management tasks, such as create, delete, resize, and upgrade clusters through this tab.
- Import and upload the following two types of OVA files that are necessary for VMware Cloud Director Container Service Extension configuration:
  - VMware Cloud Director Container Service Extension server OVA file
  - Kubernetes Template OVAs

  For information on downloading the appropriate OVA files, see Download OVA Files.
- Allow tenant users to deploy fully functional Kubernetes clusters as self-contained vApps.

Before you begin

Before you install, configure, and upgrade VMware Cloud Director Container Service Extension, you must prepare your environment. The following list outlines the prerequisites that must be in place for VMware Cloud Director Container Service Extension to operate.

- Ensure the organization virtual data centers that host tenant clusters have outbound internet connectivity.
- Ensure that you preconfigure a routed organization virtual data center network to allow you to create Kubernetes clusters. For more information, see Add a Routed Organization Virtual Data Center Network.
- Ensure VMware Cloud Director Container Service Extension server can access the public VMware Cloud Director API endpoint.
- Ensure that all required ports are open for VMware Cloud Director Container Service Extension to operate. See VMware Ports and Protocols.
- Ensure that you pre-configure VMware NSX® Advanced Load Balancer™, VMware NSX®Cloud, and NSX Advanced Load BalancerService Engine Group as outlined in Managing NSX Advanced Load Balancing. This is required to support the LoadBalancers that are deployed by the CPI of VMware Cloud Director. For more information, see Kubernetes Cloud Provider for VMware Cloud Director.
- Ensure that you use Independent Shared Named Disks in VMware Cloud Director to support Tanzu Kubernetes Grid cluster operations. For more information on named disks, see Creating and Using Named Disks. This feature is used by Container Storage Interface driver for VMware Cloud Director Named Independent Disks.

Compatibility

VMware Cloud Director Container Service Extension interacts with several other products. To ensure VMware Cloud Director Container Service Extension operates successfully, it is beneficial to examine these compatibility details.

You can check the VMware Cloud Director Container Service Extension version compatibility and interoperability on the VMware Product Interoperability Matrix.

NSX and NSX Advanced Load Balancer Compatibility

All versions of VMware Cloud Director Container Service Extension from 4.0 and above, interoperate with all versions of VMware NSX® and VMware NSX® Advanced Load Balancer™ that are interoperable with the VMware Cloud Director version used with VMware Cloud Director Container Service Extension. The interoperability matrix no longer provides compatibility information between VMware Cloud Director Container Service Extension and NSX.

Tanzu Kubernetes Grid and Kubernetes Compatibility

VMware Cloud Director Container Service Extension 4.2 supports the following Tanzu Kubernetes Grid and Kubernetes versions:

- Tanzu Kubernetes Grid 2.1.1 with Kubernetes 1.22, 1.23, and 1.24.
- Tanzu Kubernetes Grid 2.2 with Kubernetes 1.23, 1.24, and 1.25.
- Tanzu Kubernetes Grid 2.3.1 with Kubernetes 1.24, 1.25, and 1.26.
- Tanzu Kubernetes Grid 2.4 with Kubernetes 1.25, 1.26, and 1.27.
- Tanzu Kubernetes Grid 2.5 with Kubernetes 1.26, 1.27, and 1.28.

  NOTE
  Tanzu Kubernetes Grid 2.5 is supported with VMware Cloud Director Container Service Extension 4.2.1 and later.

NOTE
VMware no longer supports Tanzu Kubernetes Grid versions 1.4.3, 1.5.4 and 1.6.1. Ensure you upgrade clusters, or advise tenant users to upgrade their clusters, to use Kubernetes versions that are compatible with Tanzu Kubernetes Grid versions 2.1.1, 2.2, 2.3.1, 2.4 and 2.5. New cluster creations using Tanzu Kubernetes Grid versions 1.4.3, 1.5.4 and 1.6.1 fail with VMware Cloud Director Container Service Extension 4.1.x and 4.2.x.

Kubernetes Resources Compatibility

The following table displays the interoperability between VMware Cloud Director Container Service Extension 4.2.x with Kubernetes resources.

Kubernetes Resource	Supported Versions for 4.2	Supported Versions for 4.2.1	Supported Versions for 4.2.2	Supported Versions for 4.2.3	Documentation
Kubernetes Cloud Provider for VMware Cloud Director™	1.5.0	1.6.0	1.6.0	1.6.1	Kubernetes Cloud Provider for VMware Cloud Director Documentation
Kubernetes Container Storage Interface Driver for VMware Cloud Director™	1.5.0	1.6.0	1.6.0	1.6.0	Kubernetes Container Storage Interface driver for VMware Cloud Director
Kubernetes Cluster API Provider for VMware Cloud Director™	1.2.0	1.3.0	1.3.0	1.3.2	Kubernetes Cluster API Provider for VMware Cloud Director
RDE Projector	0.7.0	0.7.0	0.7.1	0.7.1	Not applicable

NOTE
For Node Health Check to function successfully, VMware Cloud Director Container Service Extension 4.2 requires Kubernetes Cluster API Provider for VMware Cloud Director v1.2.0. To faciliate Kubernetes Cluster API Provider for VMware Cloud Director v1.2.0 support, it is necessary to use Kubernetes Cloud Provider for VMware Cloud Director 1.5.

Software Installation

To begin configuring and using VMware Cloud Director Container Service Extension as a service provider administrator, you must install the following software.

User type	Software	Function
Service Provider	Kubernetes Container Clusters plug-in for VMware Cloud Director	Service providers can use the **CSE Management** tab of the Kubernetes Container Clusters plug-in to configure VMware Cloud Director Container Service Extension. Service providers can publish Kubernetes Container Clusters plug-in to tenants of their choice. Tenant users can use the Kubernetes Container Clusters plug-in to create Kubernetes clusters. It is necessary to download the latest version of the Kubernetes Container Clusters plug-in from the download page and upload it to VMware Cloud Director.
	VMware Cloud Director Container Service Extension server OVA file	Use for VMware Cloud Director Container Service Extension instantiation in VMware Cloud Director. For more information, see Download OVA Files.
	Kubernetes Template OVA files	Use these files to share with tenant organizations for Tanzu Kubernetes Grid cluster creation. For more information, see Download Tanzu Kubernetes Grid Templates.

Set up a Local Container Registry in an Air-gapped Environment

VMware Cloud Director Container Service Extension 4.1 and later support air-gapped environments and the use of local container registries.

> **NOTE**
> Only VMware Cloud Director Container Service Extension 4.1.x and 4.2.x provisioned clusters can operate successfully in air-gapped environments. VMware Cloud Director Container Service Extension 4.1.x and 4.2.x do not support transforming pre-existing clusters from VMware Cloud Director Container Service Extension 4.0.x into air-gapped clusters.

The following diagram details the VMware Cloud Director Container Service Extension air-gapped environment with a local container registry and the service provider workflow. This

is the default structure in VMware Cloud Director Container Service Extension 4.1.x and

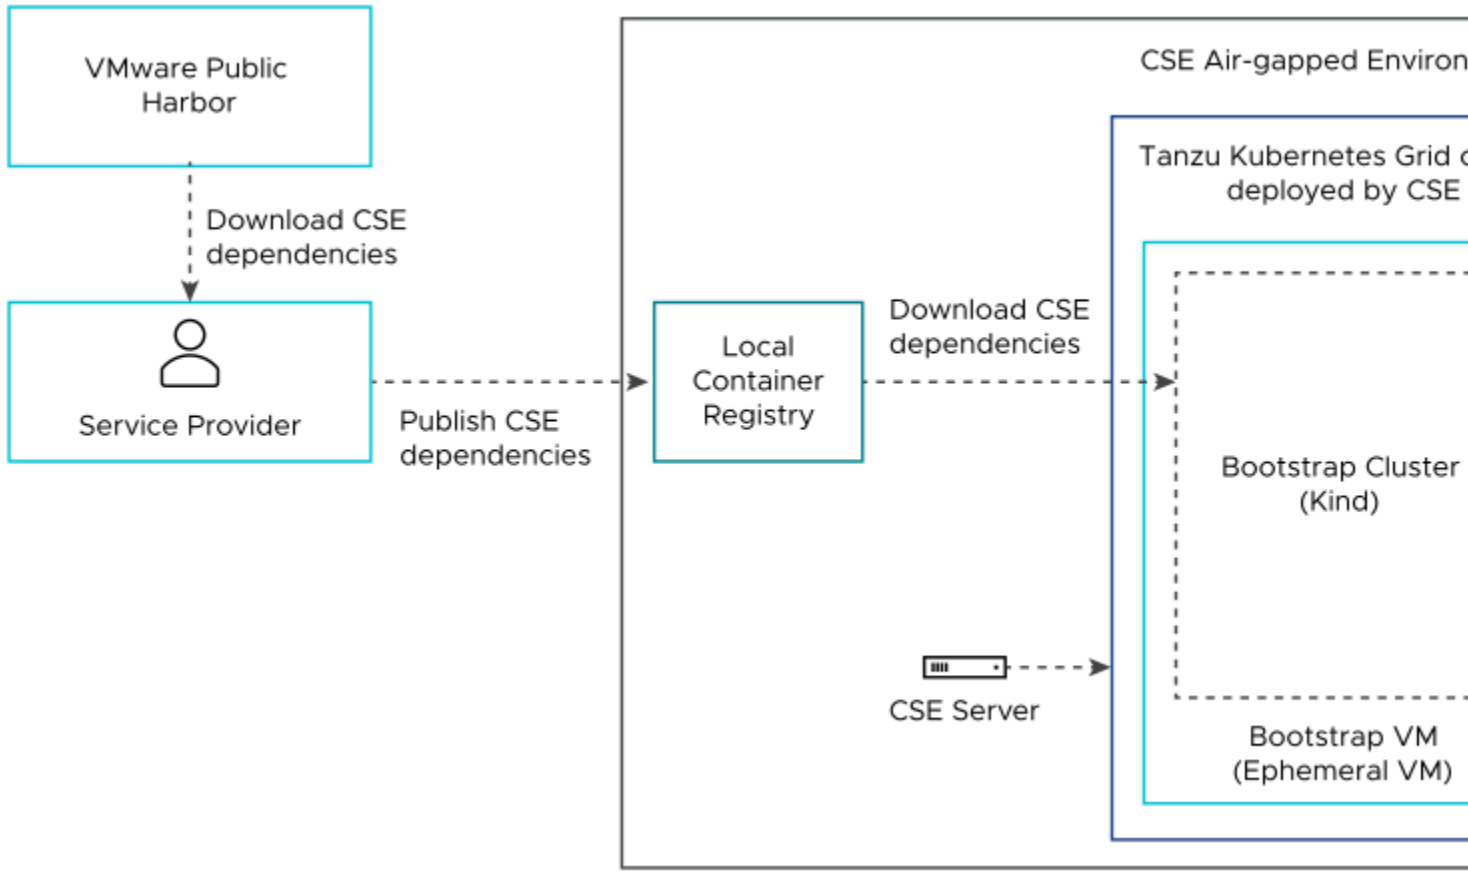

4.2.x.

Local Container Registry

Service providers set up system wide local container registries using Harbor that contain image files from VMware Public Harbor, that allow tenant users to operate VMware Cloud Director Container Service Extension in an air-gapped environment. This type of environment means tenant users can operate VMware Cloud Director Container Service Extension fully within the boundaries of their organization data center without the requirement of Internet access.

During the VMware Cloud Director Container Service Extension server configuration workflow, service providers specify the local container registry URL in the **Container Registry Settings** tab of the **Configure Settings for CSE Server** section in Kubernetes Container Clusters UI plug-in. Therefore, when tenant users attempt to create a Tanzu Kubernetes Grid cluster in the Kubernetes Container Clusters UI plug-in, it instructs Kubernetes Cluster API Provider for VMware Cloud Director, the Bootstrap VM, control planes, and worker nodes within the cluster to use the specified local container registry. For more information on configuring the local container registry details, see Server Details. Service providers can also tailor private registries for organizations that allow organizations to control what files are published and accessible to their tenant users.

Create an Air-gapped Environment

As as a service provider, you must perform the necessary steps to set up a VMware Cloud Director Container Service Extension air-gapped environment, and create a local container registry for tenant users that contains the dependencies that VMware Cloud Director Container Service Extension requires to operate.

- Ensure you have Docker engine or cli, and `imgpkg` installed on your local machine.
  - To install Docker, see Docker Desktop
  - To install `imgpkg` , see Carvel: imgpkg
- Ensure the VMware Cloud Director Container Service Extension server, and the desired tenant data centers have access to the VMware Cloud Director public end point. This is necessary as the Bootstrap VM, Kubernetes Cloud Provider for VMware Cloud Director, Kubernetes Cluster API Provider for VMware Cloud Director, and Kubernetes Container Storage Interface driver for VMware Cloud Director make calls to it.
- DNS server to resolve registry and VMware Cloud Director end point.
- Provide a Local Container Registry URL for the Bootstrap VM and all other control plane or worker VM can pull images from. For more information, see Local Container Registry.

- If the private registry is using self-signed certificates, update the Bootstrap VM and the Cluster Certificate section with the self-signed certificates. This step allows cluster virtual machines like Bootstrap VM and node VMs to trust the private registry. For more information, see Server Details.
  The following screenshot details the workflow to configure the container registry settings and certificates in Kubernetes Container Clusters plug-in for VMware Cloud Director:

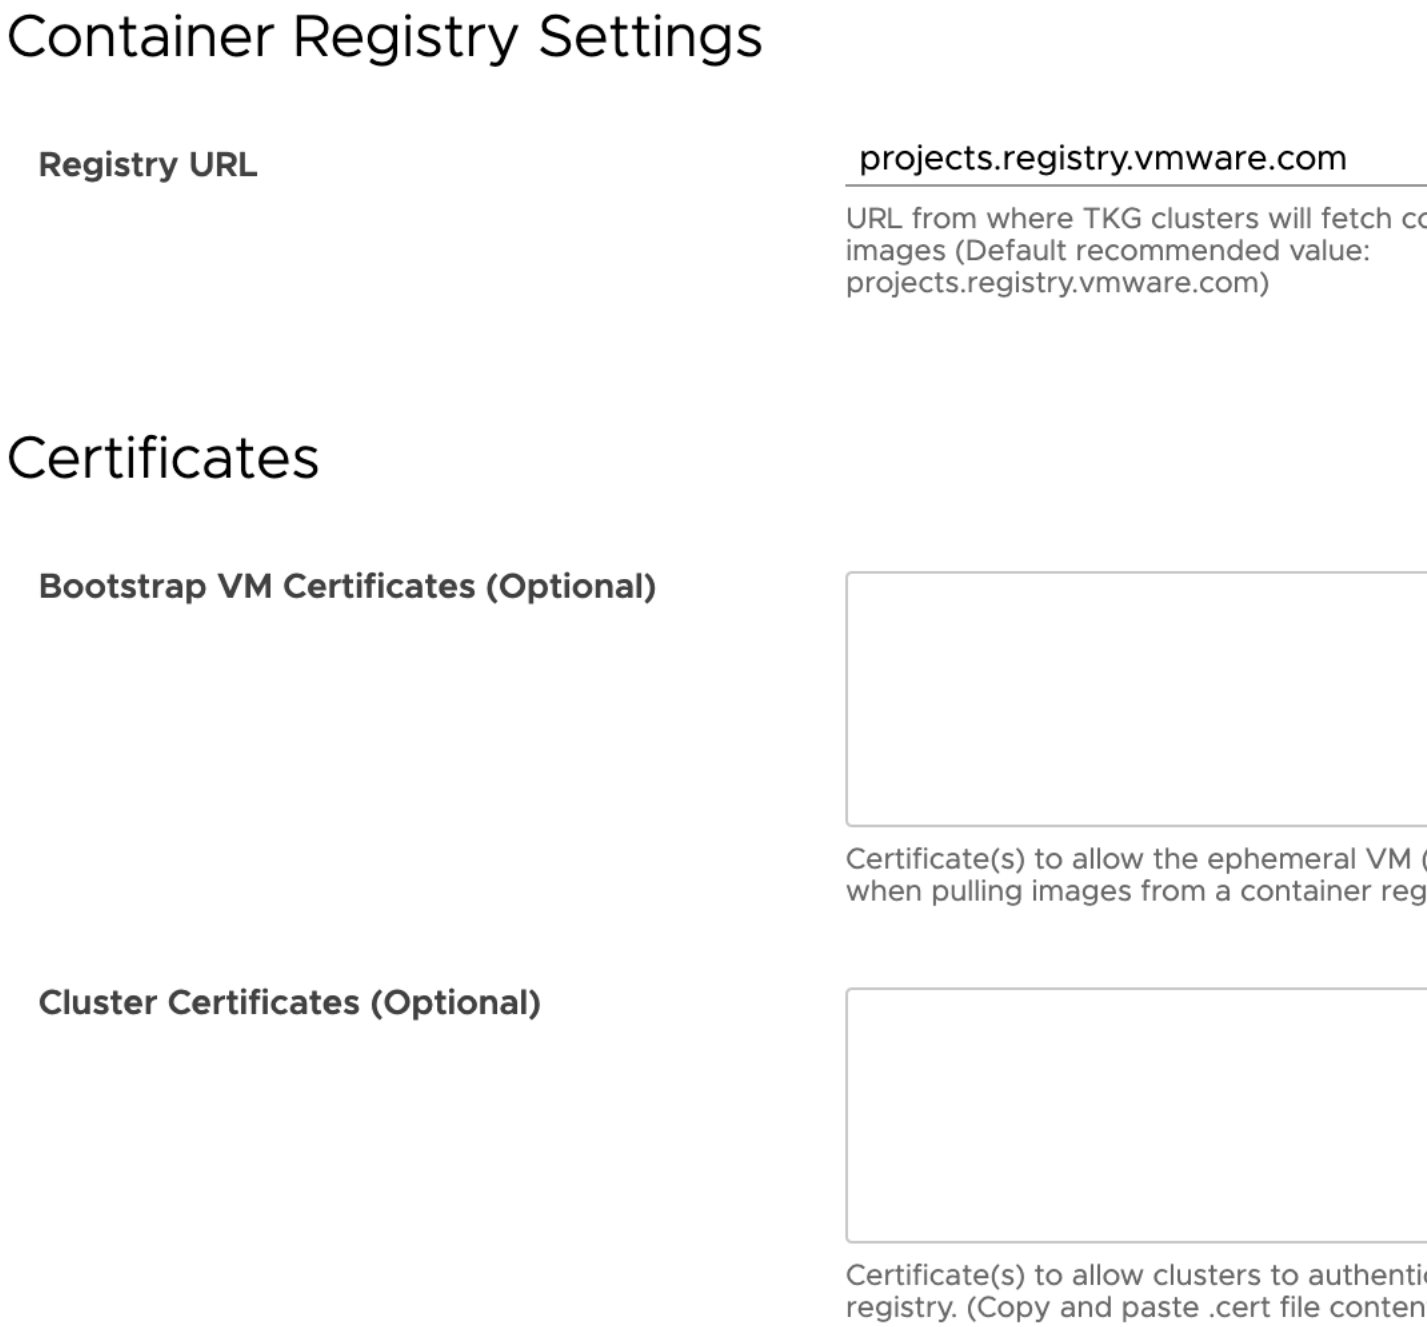

- Ensure you have access to the Internet to download the VMware Cloud Director Container Service Extension binaries from http://projects.packages.broadcom.com/ to your local machine.

1. Open Docker, and run the following command to pull the `getting-started_airgapped` image from http://projects.packages.broadcom.com.
   - For 4.2:
     ```
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-air-gapped:v0.1.2
     ```
   - For 4.2.1:
     ```
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-air-gapped:v0.1.3
     ```
   - For 4.2.2:
     ```
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-air-gapped:v0.1.4
     ```
   - For 4.2.3:
     ```
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-air-gapped:v0.1.5
     ```
2. Expand the image to gain access to the scripts and create a local directory.
   - For 4.2:
     ```
     docker create --name "temp_container" projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-airgapped:v0.1.2
     docker export "temp_container" -o temp_container.tar
     docker container rm "temp_container"
     mkdir -p temp_container_fs
     tar xvf temp_container.tar -C temp_container_fs
     cd ./temp_container_fs/src/artifact
     ```
   - For 4.2.1:
     ```
     docker create --name "temp_container" projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-airgapped:v0.1.3
     docker export "temp_container" -o temp_container.tar
     docker container rm "temp_container"
     mkdir -p temp_container_fs
     tar xvf temp_container.tar -C temp_container_fs
     cd ./temp_container_fs/src/artifact
     ```
   - For 4.2.2:
     ```
     docker create --name "temp_container" projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-airgapped:v0.1.4
     docker export "temp_container" -o temp_container.tar
     docker container rm "temp_container"
     mkdir -p temp_container_fs
     tar xvf temp_container.tar -C temp_container_fs
     cd ./temp_container_fs/src/artifact
     ```
   - For 4.2.3:
     ```
     docker create --name "temp_container" projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/getting-started-airgapped:v0.1.5
     docker export "temp_container" -o temp_container.tar
     docker container rm "temp_container"
     mkdir -p temp_container_fs
     tar xvf temp_container.tar -C temp_container_fs
     cd ./temp_container_fs/src/artifact
     ```

3. In this directory, run the following script to publish the dependencies to an organization's local container registry. There is a strict dependency on the path of the registry as displayed below.

   ```
   ./getting_started.sh "dependencies.txt" base.myregistry.company.com
   ```

   NOTE

   - The default `dependencies.txt` inside the `getting-started-airgapped` image contains the complete list of containers that are necessary to use VMware Cloud Director Container Service Extension in an air-gapped environment for the supported versions of Tanzu Kubernetes Grid. For supported versions of Tanzu Kubernetes Grid, see the Compatibility section.
   - You can edit the `dependencies.txt` file to customize what components you want to download. You can supress the processing of an air-gapped component by inserting a `#` in front of the corresponding line in `dependencies.txt` .

4. It is necessary to upload Tanzu Core Packages from projects.packages.broadcom.com to the local container registry. To do this, use the instructions from the appropriate Tanzu Kubernetes Grid release below.

   - Tanzu Kubernetes Grid 2.4
   - Tanzu Kubernetes Grid 2.5

     NOTE
     Tanzu Kubernetes Grid 2.5 is available with VMware Cloud Director Container Service Extension 4.2.1 and later.

5. If Tanzu Kubernetes Grid 2.3.1, 2.4, and 2.5 clusters need to be created, it is necessary to upload the Tanzu Kubernetes Grid 2.3.1, 2.4, and 2.5 CLI plugin to the registry path `<registry_url>/tanzu_cli/plugins/plugin-inventory:latest` . For detailed instructions, see Move the Tanzu CLI and CLI Plugins to Your Internet-Restricted Environment.

   NOTE
   With VMware Cloud Director Container Service Extension 4.2 and later, there is a strict dependency on the path where the Tanzu Kubernetes Grid images are uploaded.

   - If the registry base is `base.myregistry.company.com` , the Tanzu Kubernetes Grid content should be uploaded into `base.myregistry.company.com/tkg` .
   - If the registry base is `base.myregistry.company.com` , the VMware Cloud Director images should be uploaded into `base.myregistry.company.com/vmware-cloud-director` .
   - If the registry base is `base.myregistry.company.com` , the Tanzu CLI plugin for TKG 2.3.1, TKG 2.4, and TKG 2.5 should be uploaded into `base.myregistry.company.com/tanzu_cli/plugins/plugin-inventory:latest` .

   This implies that the sub-paths `tkg` , `vmware-cloud-director` and `tanzu_cli` must reside next to eachother.

An air-gapped environment is now set up so the organization can use VMware Cloud Director Container Service Extension without the requirement of Internet access with all the required resources in a local container registry.

Kubernetes Container Clusters Plug-in for VMware Cloud Director

The Kubernetes Container Clusters plug-in is the graphical user interface of the VMware Cloud Director Container Service Extension for VMware Cloud Director. Use this section to learn how to use this plug-in to perform and manage vital operations in VMware Cloud Director Container Service Extension.

With VMware Cloud Director Container Service Extension 4.1 and later, service providers can use the **CSE Management** tab in the Kubernetes Container Clusters plug-in to configure the VMware Cloud Director Container Service Extension server, in addition to creating Tanzu Kubernetes Grid clusters, and performing VMware Cloud Director Container Service Extension management tasks.

You can download the latest version of the Kubernetes Container Clusters plug-in and upload it to VMware Cloud Director. For upload instructions, see Upload a Plug-in.

To check the compatibility of VMware Cloud Director and Kubernetes Container Clusters versions, see the Product Interoperability Matrix.

To allow tenants to create Kubernetes clusters, you must publish the Kubernetes Container Clusters plug-in to those organizations. For more information, see Publish or Unpublish a Plug-in from an Organization.

> **NOTE**
> If you have previously used the Kubernetes Container Clusters plug-in with VMware Cloud Director, you must deactivate the plug-in before activating a newer version, as only one version of the plug-in can operate at one time in VMware Cloud Director. After you activate a new plug-in, to begin using it, you must refresh your Internet browser. For more information, see Activate or Deactivate a Plug-in. For more information about managing VMware Cloud Director plug-ins, see Managing Plug-Ins.

Set up the VMware Cloud Director Container Service Extension server through the Kubernetes Container Clusters UI plug-in

You can set up the VMware Cloud Director Container Service Extension server through the **CSE Management** tab in Kubernetes Container Clusters UI plug-in. You can perform the task in this tab to ensure Tanzu Kubernetes Grid clusters, and relevant workloads operate successfully.

The **CSE Management** tab in Kubernetes Container Clusters UI plug-in has the following sections:

- **Getting Started**
- **Guidelines**
- **Server Details**

Getting Started

The **Getting Started** section is a landing page for VMware Cloud Director Container Service Extension management. Learn how to set up VMware Cloud Director Container Service Extension in VMware Cloud Director through the Kubernetes Container Clusters UI plug-in, to allow tenant users to create Kubernetes clusters.

There are different sections in VMware Cloud Director where you can perform the necessary tasks to create a suitable environment for VMware Cloud Director Container Service Extension to operate successfully in VMware Cloud Director.

Table 1: Getting Started Page Contents

Sections	Description
Download OVAs	This section links to the locations where you can download the following two types of OVA files that are necessary for VMware Cloud Director Container Service Extension configuration: • VMware Cloud Director Container Service Extension server OVA file • Kubernetes template OVA files For information on downloading the appropriate OVA files, see Download OVA Files.
Create Catalogs and Upload OVAs	This section links to the Catalogs section in VMware Cloud Director where you can perform the following actions: • Leverage catalogs to store and retrieve VMware Cloud Director Container Service Extension server OVA files, and maintain a repository of Kubernetes Template OVAs. • Create a catalog in a provider-managed organization, and upload VMware Cloud Director Container Service ExtensionServer OVA file for easy access. • Create a shared catalog in a provider-managed organization and upload Kub ernetes Template OVAs. For more information, see Create Catalogs and Upload OVA Files.
Setting up the Configuration for CSE Server	This section initiates the VMware Cloud Director Container Service Extension server configuration process. For more information, see Configure the VMware Cloud Director Container Service Extension server.
Add VM Sizing Policies to Organization VDCs	This section links to the **Organization VDCs** section in VMware Cloud Director, where you can assign VM sizing policies to organization VDCs. For more information, see Add Tanzu Kubernetes Grid VM Sizing Policies to Organization Virtual Data Centers.
Create a User with CSE Admin Role	This section links to the **Users** section in VMware Cloud Director, where you can create a user with the **CSE Admin Role** role. This role grants administration privileges to the user for VMware Cloud Director Container Service Extension administrative purposes. For more information, see Create a User with CSE Admin Role.
Start CSE Server	This section links to the **vApps** section in VMware Cloud Director where you can create a vApp from the uploaded VMware Cloud Director Container Service Extension server OVA file to start the VMware Cloud Director Container Service Extension server. For more information, see Create a vApp from VMware Cloud Director Container Service Extension server OVA file.

Guidelines

The **Guidelines** section in the **CSE Management** tab details the cloud resources that the **Getting Started** workflow automatically creates.

Table 2: Guideline page contents

Section	Description
Rights & Roles	This section details the rights bundle and user roles that that the **Getting Started** automatically creates.
TKG VM Sizing Policies	It is necessary to manually add the VM sizing policies to organization virtual data centers. K ubernetes clusters that do not have one of these VM sizing policies can experience resource limit errors. This section details the Tanzu Kubernetes Grid VM Sizing Policies you can add to organization virtual data centers. For information on how to perform this task, see Add Tanzu Kubernetes Grid VM Sizing Policies to Organization Virtual Data Centers.

Server Details

The **Server Details** section in the Kubernetes Container Clusters UI plug-in details the VMware Cloud Director Container Service Extension server configuration details.

You can use this section to view the current server configuration details, and if you want to update these settings through the Kubernetes Container Clusters UI plug-in, click **Update Server**. For more information, see Update the VMware Cloud Director Container Service Extension Server. For more information on server configuration details, see Configure the VMware Cloud Director Container Service Extension Server Settings.

Node Health Check Configuration

You can configure, activate, and deactivate the Node Health Check parameters in Tanzu Kubernetes Grid clusters through the Kubernetes Container Clusters UI plug-in.

The Node Health Check feature comprises of two parts:

- Detection
- Remediation

> **NOTE**
> Node Health Check and Auto Repair on Errors are different in functionality. Node Health Check detects and remediates unhealthy nodes in the cluster only after the cluster goes to an **Available** status, while Auto repair on errors reattempts cluster creation if cluster goes to error state before cluster status becomes **Available**.

> **NOTE**
> Node Health Check is deactivated by default in VMware Cloud Director Container Service Extension 4.1 and newer versions.

Node Failure Detection

VMware Cloud Director Container Service Extension 4.1 and newer versions can detect when a node in a Tanzu Kubernetes Grid cluster becomes unhealthy. When a node is in an unhealthy state, the Kubernetes Container Clusters UI plug-in reflects the available and desired node count in the cluster information page, and also the failure appears in the **Events** section of the same page.

A node can become unhealthy for the following reasons but not limited to

- Network outages
- Power interruptions
- Low node speed due to high memory, CPU or disk utilization
- Node startup failure

- Failure to join the cluster

Node Remediation

From VMware Cloud Director Container Service Extension 4.1, the Node Health Check feature detects node failure in Tanzu Kubernetes Grid clusters, and automatically replaces unhealthy Kubernetes nodes with new nodes. The Node Health Check parameters are required global settings for the VMware Cloud Director Container Service Extension server setup, and server update workflows, which are used by Kubernetes Container Clusters UI plug-in to create clusters, or update settings for clusters in all organizations. For more information, see Update the VMware Cloud Director Container Service Extension Server. Service providers can return to the **Update Server** tab at any time to reconfigure Node Health Check parameters. If service providers do not specifically configure the Node Health Check parameters, the following default values are set:

Table 3: Node Health Check Configuration

Node Health Check Parameter	Default Value	Description
Max Unhealthy Nodes	100%	Remediation is suspended when the percentage of unhealthy nodes exceeds this value. When the default value is 100%, this means the cluster is always remediated. When the default value is 0%, this means the cluster does not remediate.
Node Startup Timeout	900 seconds	If a node does not start in this time frame, it is considered unhealthy and is remediated. For a given VMware Cloud Director environment, it is recommended for service providers to set Node Health Check parameter to be at least twice the time for a VM to be created and bootstrapped.
Node Status "Not Ready" Timeout	300 seconds	If a newly joined node cannot host workloads for longer than this timeout, it is considered unhealthy and is remediated.
Node Status "Unknown" Timeout	300 seconds	If a healthy node is unreachable for longer than this timeout, it is considered unhealthy and is remediated.

Tenant users use the Node Health Check parameters set by the service provider for their organization when they create clusters. For more information, see Create a Tanzu Kubernetes Grid Cluster.

> **NOTE**
> When service providers update the Node Health Check parameters, the existing Node Health Check parameters on the Tanzu Kubernetes Grid clusters that are already deployed are not modified.

Activate or Deactivate Node Health Check in a VMware Cloud Director Container Service Extension 4.0.x Cluster

Tenant users can also activate or deactivate Node Health Check on clusters that were created in VMware Cloud Director Container Service Extension 4.0.x.

The following steps outline how tenant users can perform this action:

1. Log in to VMware Cloud Director portal, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Click the cluster name, and in the cluster information page, click **Settings**.
3. Activate or deactivate **Node Health Check** toggle, and click **Save**.

View Clusters

You can use the Kubernetes Container Clusters UI plug-in of VMware Cloud Director to view clusters and cluster details.

To view your clusters in the Kubernetes Container Clusters UI plug-in, from the VMware Cloud Director navigation menu, select **More** > **Kubernetes Container Clusters**. A list of Kubernetes clusters created by VMware Cloud Director Container Service Extension displays, along with some basic cluster information.

Depending on your role, Kubernetes Container Clusters UI plug-in has two role-specific views:

Table 4: Kubernetes Container Clusters UI plug-in user views

User Type	View
Service provider	Service providers can view all clusters in all organizations.
Tenant user	Tenant users can only view clusters in their organization that they have visibility for.

Manage Clusters

You can manage your clusters in the Kubernetes Container Clusters UI plug-in. Learn about different management functions you can perform during the lifecycle of a cluster.

To...	Do this...
Create a cluster	Click **New** on the top left to open the cluster creation wizard. For more information on how to create a cluster, refer to Create a VMware Tanzu Kubernetes Grid Cluster.
Download Kube Config	Click **Download Kube Config** on the top left. For more information on Kube Config file, refer to the Kubernetes website.
Upgrade a cluster	Select a cluster in the datagrid, and click **Upgrade** on the top left. For more information, see Upgrade a Tanzu Kubernetes Grid Cluster.
Attach to Tanzu Mission Control	Click **Attach to TMC** if you want to attach the cluster to Tanzu Mission Control. For more information, see Attach a Cluster.
Detach from Tanzu Mission Control	Click **Detach from TMC** if you want to detach the cluster from Tanzu Mission Control. For more information, see Detach a Cluster.
Manage on Tanzu Mission Control	Click **Manage on TMC** if you want to manage the cluster on Tanzu Mission Control. For more information, see Manage a Cluster.
Delete a cluster	Select a cluster in the datagrid, and click **Delete** on the top left. For more information, see Delete a Kubernetes Cluster.
Resize a cluster	For information on how to resize a cluster, see Resize a Tanzu Kubernetes Grid Cluster.

VMware Cloud Director Container Service Extension Server

Ensure the VMware Cloud Director Container Service Extension prerequisites are in place, and learn directions on how to install VMware Cloud Director Container Service Extension server.

Use the Kubernetes Container Clusters plug-in in VMware Cloud Director to navigate through the steps in this section. The Kubernetes Container Clusters UI plug-in outlines the configuration workflow of the VMware Cloud Director Container Service Extension server, and directs you to the VMware Cloud Director UI to perform configuration steps. It is important to follow the steps in sequential order, and examine the prerequisites in each section to ensure the appropriate conditions are in place before you can perform each task.

Co-existence of VMware Cloud Director Container Service Extension servers with Kubernetes Container Clusters UI plug-in 4.x

VMware Cloud Director Container Service Extension 3.1.x and VMware Cloud Director Container Service Extension 4.x are two different technical stacks. It is not possible to upgrade between the two versions. To serve VMware Cloud Director Container Service Extension 3.1.x cluster requirements, a server co-existence model exists and is supported through the Kubernetes Container Clusters UI plug-in 4.x.

If tenant users have Tanzu Kubernetes Grid clusters that were created in VMware Cloud Director Container Service Extension 3.1.x and they want to adopt the Tanzu Kubernetes Grid feature set of VMware Cloud Director Container Service Extension 4.x, it is necessary to install VMware Cloud Director Container Service Extension 4.x server. If a tenant user wants to keep VMware Cloud Director Container Service Extension 3.1.x clusters, both Tanzu Kubernetes Grid and native, they can use a co-existence model in the same VMware Cloud Director environment, where the servers exist side by side.

When a tenant user wants to create a new cluster, they can only create Tanzu Kubernetes Grid clusters, and can only create them in VMware Cloud Director Container Service Extension 4.x through the Kubernetes Container Clusters UI plug-in 4.x.

The aim of the server co-existence model is that eventually tenant users will migrate Kubernetes workloads from VMware Cloud Director Container Service Extension 3.1.x Tanzu Kubernetes Grid clusters to VMware Cloud Director Container Service Extension 4.x Tanzu Kubernetes Grid, and remove VMware Cloud Director Container Service Extension 3.1.x from VMware Cloud Director.

Frequently Asked Questions

Table 5: Frequently Asked Questions

Question	Answer
Do I need the VMware Cloud Director Container Service Extension 3.1.x server?	It necessary to keep the VMware Cloud Director Container Service Extension 3.1.x only if you want to continue to manage VMware Cloud Director Container Service Extension 3.1.x clusters
Do I need two Kubernetes Container Clusters UI plug-ins to support the co-existence model?	No. You can manage both VMware Cloud Director Container Service Extension 3.1.x and VMware Cloud Director Container Service Extension 4.x clusters through the Kubernetes Container Clusters UI plug-in 4.x.
When can I remove VMware Cloud Director Container Service Extension 3.1.x server and keep only the VMware Cloud Director Container Service Extension 4.x server?	Once you delete all clusters managed by VMware Cloud Director Container Service Extension 3.1.x, you can unregister and remove VMware Cloud Director Container Service Extension 3.1.x server
Can I migrate VMware Cloud Director Container Service Extension 3.1.x clusters and manage them through the VMware Cloud Director Container Service Extension 4.x server?	No. It is necessary to create a new VMware Cloud Director Container Service Extension 4.x cluster, manually migrate apps over from VMware Cloud Director Container Service Extension 3.1.x cluster to VMware Cloud Director Container Service Extension 4.x cluster, and then delete VMware Cloud Director Container Service Extension 3.1.x cluster.

Question	Answer
Can I migrate native Kubernetes clusters from VMware Cloud Director Container Service Extension 3.1.x server to VMware Cloud Director Container Service Extension 4.x server?	No, VMware Cloud Director Container Service Extension 4.x server does not support native Kubernetes clusters. To support this type of cluster, it is necessary to use the co-existence server model with the support of Kubernetes Container Clusters UI plug-in 4.x.

VMware Cloud Director Container Service Extension Server High Availability

From VMware Cloud Director Container Service Extension 4.0, the VMware Cloud Director Container Service Extension server is highly available. The high availability features display how you can use the VMware Cloud Director Container Service Extension server to optimize workload performance.

VMware Cloud Director Container Service Extension server high availability means it can remain continuously functional, even when issues occur. If an incident occurs, the highly available server can shift workloads and configurations away from the affected nodes. The capability to minimize unplanned interruptions allows the server to be operational at all times.

The following table details the features of VMware Cloud Director Container Service Extension server high availability:

Features	Description
Horizontal scalability of the server	You can run VMware Cloud Director Container Service Extension service in multiple virtual machines simultaneously.
Elimination of redundant software and applications	Ability of the software to react to failures. Systemctl restarts VMware Cloud Director Container Service Extension service whenever it fails.
No single point of failure	A failure in a single component does not crash the entire infrastructure. Even when the VMware Cloud Director Container Service Extension server fails, VMware Cloud Director can accept requests. Cluster upgrade and resize operations do not require the VMware Cloud Director Container Service Extension server to be running.
Fault tolerance	The system can recover from failures. VMware Cloud Director Container Service Extension provides Auto-repair as a feature in the cluster creation workflow. For more information, see Create a Tanzu Kubernetes Grid Cluster.
Disaster recovery	Recovery from a catastrophic event where a physical data center or other infrastructure is damaged. As VMware Cloud Director Container Service Extension is stateless, all of its configuration information is stored in the form of an RDE in the VMware Cloud Director database. All the cluster RDEs are also stored in the VMware Cloud Director database. RDEs can be backed up along with OVDC using VMware Cloud Director Disaster Recovery strategies. To prepare for a cluster failure, it is recommended to use VMware Cloud Director™ Object Storage Extension™ to back up cluster information. For more information, see Backing up and Restoring Kubernetes Clusters.

VMware Cloud Director Container Service Extension Server Prerequisites

To prepare an environment that allows VMware Cloud Director Container Service Extension server to function, it is necessary to ensure you have all the prerequisite steps in place.

VMware Cloud Director Setup Prerequisites

It is necessary to configure required components before you can install the VMware Cloud Director Container Service Extension server successfully on VMware Cloud Director. You can use existing resources from your VMware Cloud Director or create new parameters.

Table 6:

Required component	Description
An organization	Before you can configure VMware Cloud Director Container Service Extension server, it is necessary to create an organization in VMware Cloud Director. The organization is considered a Cloud Provider managed organization that hosts VMware Cloud Director Container Service Extension server. For more information, see Create an Organization.
A virtual data center (VDC) within the organization	Ability to host VMware Cloud Director Container Service Extension server vApp. For more information, see Create an Organization Virtual Data Center.
Network connectivity	Network connectivity between the machine where VMware Cloud Director Container Service Extension is installed, and the VMware Cloud Director server. VMware Cloud Director Container Service Extension communicates with VMware Cloud Director using VMware Cloud Director public API endpoint.

Download OVA Files

Before you begin to configure VMware Cloud Director Container Service Extension server and deploy Kubernetes clusters to tenant organizations, you must download the VMware Cloud Director Container Service Extension OVA file and Tanzu Kubernetes Grid template OVA files.

OVA File	Download location
VMware Cloud Director Container Service Extension server OVA file	VMware Cloud Director Container Service Extension
Kubernetes template OVA files	Tanzu Kubernetes Grid Templates **NOTE** Ubuntu 2004 Kubernetes Non-FIPS Templates only

Create Catalogs and Upload OVA Files

Learn how to create a catalog in VMware Cloud Director, and upload VMware Cloud Director Container Service Extension OVA files that you downloaded into catalogs. You can use VMware Cloud Director Container Service Extension server file for VMware Cloud Director Container Service Extension instantiation, and the Kubernetes Template OVA file to share with tenant organizations for Tanzu Kubernetes Grid cluster creation.

Ensure you download the following two OVA files:

- VMware Cloud Director Container Service Extension Server OVA file
- Kubernetes Template OVA file

For more information, see Download OVA Files.

1. Log in to VMware Cloud Director and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Click **CSE Management** > **Create Catalogs and Upload OVAs** > **Go to Catalogs**.
3. In VMware Cloud Director UI, select an organization, and click **Go to Catalogs**.
4. Select **New**, enter the name and, optionally, a description of the catalog.

   NOTE
   It is recommended to have separate catalogs for each type of OVA file.
5. Click **OK**.
6. To upload OVA files, select a catalog, and from the left column, select **vApp Templates** > **New**.
7. Click the **Upload** icon to browse to a location accessible from your computer, and select the OVA file.
8. Review the details of the OVA file and click **Next**.
9. Enter a name and, optionally, a description, and click **Next**.
10. Click **Finish**.

The OVA files appear in a new catalog in VMware Cloud Director.

Sharing Tanzu Kubernetes Grid Templates.

Add Tanzu Kubernetes Grid VM Sizing Policies to Organization Virtual Data Centers

To avoid resource limit errors in clusters, it is necessary to add Tanzu Kubernetes Grid VM sizing policies to organization virtual data centers. The Tanzu Kubernetes Grid VM sizing policies are automatically created in the VMware Cloud Director Container Service Extension server configuration process and appear as a selectable option when you add Tanzu Kubernetes Grid VM sizing policies to an organization virtual data center.

A VM sizing policy defines the compute resource allocation for virtual machines within an organization VDC. The compute resource allocation includes CPU and memory allocation, reservations, limits, and shares. By using VM sizing policies, you can restrict the compute resources consumption for all virtual machines within an organization VDC to predefined sizes. Tanzu Kubernetes Grid clusters have the following sizing policies:

Table 7: Tanzu Kubernetes Grid Cluster Sizing Policies

Sizing Policy	Description	Values
TKG small	Small VM sizing policy for a Kubernetes cluster node	2 CPU, 4 GB memory
TKG medium	Medium VM sizing policy for a Kubernetes cluster node	2 CPU, 8 GB memory
TKG large	Large VM sizing policy for a Kubernetes cluster node	4 CPU, 16 GB memory

Sizing Policy	Description	Values
TKG extra-large	Extra-large VM sizing policy for a Kubernetes cluster node	8 CPU, 32 GB memory

1. Log in to VMware Cloud Director and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Click **CSE Management** > **Add VM Sizing Policies to Organization VDCs** > **Go to Organization VDCs**.
3. In VMware Cloud Director UI, select an organization VDC, and from the left panel, under **Policies**, select **VM Sizing**.
4. Click **Add**.
5. From the data grid, select the Tanzu Kubernetes Grid sizing policy you want to add to the organization, and click **OK**.
6. Click **Set as Default**.

NOTE

- To unpublish a VM sizing policy, select the sizing policy, and click **Remove**.
- To change the default Tanzu Kubernetes Grid VM sizing policy to another, repeat Step 5 and 6.

Configure the VMware Cloud Director Container Service Extension Server Settings

You must complete the VMware Cloud Director Container Service Extension server configuration process before it can operate. This workflow automatically creates a **Kubernetes Clusters** rights bundle, **CSE Admin Role** role, **Kubernetes Cluster Author** global role, and VM sizing policies. In this process, the **Kubernetes Clusters** rights bundle and **Kubernetes Cluster Author** role are automatically published to all tenants.

Before you can configure the VMware Cloud Director Container Service Extension server, ensure you have met the following prerequisites:

- Complete the rerequisites detailed in VMware Cloud Director Setup Prerequisites
- Upload the OVA files detailed in Create Catalogs and Upload OVA Files.

You can specify the following settings when you configure the VMware Cloud Director Container Service Extension server.

Table 8: Server Configuration Parameters

Server Parameter	Description
Current CSE Server Component Versions	**CAPVCD Version**: Kubernetes Cluster API Provider for VMware Cloud Director
	Kubernetes Cloud Provider for VMware Cloud Director
	Kubernetes Container Storage Interface driver for VMware Cloud Director
	Optional: Github Personal Access Token
	Optional: Bootstrap Cluster VM Sizing Policy
Current Proxy Settings: Optional	NO_PROXY: List of comma-seperated domains without spaces
	HTTP_PROXY: Address of HTTP proxy server
	HTTPS_PROXY: Address of HTTPS proxy server **NOTE** Proxy settings in VMware Cloud Director Container Service Extension are used only for image downloads in air-gapped environments or from external repositories. Proxy communication between the internal pods of a routed Organisation VDC and the VMware Cloud Director API is not supported.
Current Syslog Location: Optional	Host: Domain name

Server Parameter	Description
	Port: Port number
Current Node Health Check Settings	• Node Startup Timeout • Node "Not Ready" Timeout • Node "Unknown" Timeout • Max Unhealthy Nodes For more information, see Node Health Check Configuration.
Current Container Registry Settings	Registry URL: The URL where Tanzu Kubernetes Grid clusters fetch container images. **NOTE** To use VMware Cloud Director Container Service Extension in an air-gapped environment, it is necessary to enter a local container registry URL. For more information, see Set up a Local Container Registry in an Air-gapped Environment.
Certificates	Bootstrap VM Certificates: Certificates that allow the ephemeral VM, that is created during cluster creation, to authenticate with. For example, when pulling images from a container registry. It is necessary to copy and paste the `.cert` file contents.
	Cluster Certificates: Certificates that allow clusters to authenticate with. For example, when pulling images from a container registry. It is necessary to copy and paste the `.cert` file contents.

1. Log in to VMware Cloud Director and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Click **CSE Management** > **Setting up the Configuration for CSE Server** > **Start**.
3. In VMware Cloud Director UI, in the **Configure Settings for CSE Server** window, in the **Set Up Prerequisites** section, click **Start**.
4. Click **Next**.
5. In the **Set Configuration Parameters** section, configure the VMware Cloud Director Container Service Extension server settings. For a list of these settings, see Server Details Section.

   NOTE
   To revert to default configuration values, click **Restore Defaults**.
6. Click **Submit**.

   Server configuration entity is successfully created.

Create a User with CSE Admin Role

You can create a user with**CSE Admin Role** in the system organization. The**CSE Admin Role** allows a user to perform administrative tasks in VMware Cloud Director Container Service Extension. You can use these user credentials as OVA deployment parameters when you start the VMware Cloud Director Container Service Extension server.

- Ensure that the user has**system administrator** privileges. For more information, refer to System Administrator Rights.
- It is necessary to complete the VMware Cloud Director Container Service Extension server configuration before you can assign the**CSE Admin Role**. For more information, see Configure VMware Cloud Director Container Service Extension Server.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Click **CSE Management** > **Create a User with CSE Admin Role** > **Go to Users**.
3. In VMware Cloud Director UI, click **New**.
4. Enter a user name and an API Token for the new user. For more information, see Generate an API Access Token.

   NOTE
   When you generate an API Access Token, you must copy the token, because it appears only once. After you click **OK**, you cannot retrieve this token again, you can only revoke it. You must use the same API Access Token when you deploy the VMware Cloud Director Container Service Extension server OVA file. For more information, see Create a vApp from VMware Cloud Director Container Service Extension server OVA file.
5. Select the **Enable** toggle to activate the user upon creation.
6. From the **Available roles** drop-down menu, select **CSE Admin Role**.
7. Enter contact information for the user. You can enter the full name, email address, phone number, and instant messaging ID.
8. Set the quotas for the user.
   1. You can set a limit of the virtual machines owned by the user, or select **Unlimited**.
   2. You can set a limit of the running virtual machines owned by the user, or select **Unlimited**.
9. Click **Save**.

Start the VMware Cloud Director Container Service Extension Server

Ensure you complete necessary the steps to start the VMware Cloud Director Container Service Extension server successfully.

- Create a vApp from VMware Cloud Director Container Service Extension server OVA file.
- Configure the VMware Cloud Director Container Service Extension server vApp deployment lease
- Power on the VMware Cloud Director Container Service Extension server.

Create a vApp from VMware Cloud Director Container Service Extension server OVA file

You can create a vApp from the VMware Cloud Director Container Service Extension server OVA file. The vApp is created in the organization where you want to run the VMware Cloud Director Container Service Extension server vApp.

- Ensure that you have met the VMware Cloud Director Setup Prerequisites.
- Set the lease of the chosen organization to never expire.

1. Log in to VMware Cloud Director and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Click **CSE Management** > **Start CSE Server** > **Go to vApps**.
3. In the VMware Cloud Director Tenant Portal, select an organization, and click **Go to vApps**.
4. In VMware Cloud Director UI, click **Applications**.
5. In the **Virtual Applications** tab, click **New**, and choose one of the following options:
   - **Add vApp From OVF**: Select this option to upload the OVA file from your local machine.
     a. Select a virtual data center and click **Next**.
     b. Click the **Browse** icon, navigate to the `VMware_Cloud_Director_Container_Service_Extension-4.x.x.ova` file, and click **Open**.
     c. In the **Accept Licenses** window, select the **I agree and accept the above license agreements** check box.
     d. Enter a vApp name, a description, and click **Next**.
     e. Review the default configuration for resources, networking, custom properties, hardware, and optionally edit details where necessary.
     f. In the **Custom Properties** window, configure the following settings:

CSE Server Properties	Description
VCD host	The URL that organization members can use to access VMware Cloud Director UI.
CSE service account's username	The user name of**CSE Admin** user in the organization.
CSE service account's API Token	The API access token you generate when you create a user with **CSE Admin** role.
CSE service account's org	The organization that the user with the**CSE Admin** Role belongs to, and that the VMware Cloud Director Container Service Extensionserver deploys to.
CSE service vApp's org	The organization that the vApp deploys to.

     For more information, see Create a User with CSE Admin Role.
     g. Click **Finish**.
   - **Add vApp From Catalog**: Select this option if the OVA file is in a catalog in an organization in VMware Cloud Director.
     a. Select a virtual data center and click **Next**.
     b. Select the OVA file to import and click **Next**.
     c. In the **Accept Licenses** window, click **Accept** > **Next**.

        NOTE
        If you do not want to accept the **End User License Agreement**, click**Reject** > **Start Over**.
     d. Enter a vApp name, optionally a description, runtime lease and storage lease, and click **Next**.
     e. Review the default configuration for resources, compute policies, hardware, networking, and edit details where necessary.

f. In the **Custom Properties** window, configure the following settings:

CSE Server Properties	Description
VCD host	The URL that organization members can use to access VMware Cloud Director UI.
CSE service account's username	The user name of**CSE Admin** user in the organization.
CSE service account's API Token	The API access token you generate when you create a user with **CSE Admin** role.
CSE service account's org	The organization that the user with the**CSE Admin** role belongs to, and that the VMware Cloud Director Container Service Extension server deploys to.
CSE service vApp's org	The organization that the vApp deploys to.

For more information, see Create a User with CSE Admin Role.

g. Click **Finish**.

6. In the **Virtual Applications** tab, in the bottom left of the vApp, click **Actions** > **Power** > **Start**.

NOTE
If you want to log into the VMware Cloud Director Container Service Extension server at any point, it is necessary to set up an auto-generated password in **Guest-Customization** in VMware Cloud Director, and restart the virtual machine. For more information, see Change the Guest OS Customization of a Virtual Machine.

The vApp is created from the VMware Cloud Director Container Service Extension server OVA file.
Configure the VMware Cloud Director Container Service Extension Server vApp Deployment Lease.

Configure the VMware Cloud Director Container Service Extension Server vApp Deployment Lease

It is necessary to configure this setting to **Never Expires** to avoid the VMware Cloud Director Container Service Extension vApp from powering off, or deleting automatically due to inactivity.

- Create a provider-managed organization to store a vApp from the uploaded VMware Cloud Director Container Service Extension server OVA. For more information, see Create an Organization.
- Create a vApp from VMware Cloud Director Container Service Extension server OVA file.
- To learn more about leases, see Understanding Leases.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **Resources** > **Cloud Resources**.
2. From the left panel, select **Organizations**, and select an organization.
3. In the cluster information window, from the left panel, select **Policies**.
4. In the **vApp Leases** section, click **Edit**.
5. In the **vApp Leases** window, from the **Maximum Runtime Lease** dropdown list, select **Never Expires**.

Power on the VMware Cloud Director Container Service Extension Server.

Power on the VMware Cloud Director Container Service Extension Server

Learn how to turn on the VMware Cloud Director Container Service Extension server.

Ensure that you Create a vApp from VMware Cloud Director Container Service Extension server OVA file.

This task is the final step for service providers to perform before the VMware Cloud Director Container Service Extension server can operate.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Click **CSE Management** > **Start CSE Server** > **Go to vApps**.
3. In the VMware Cloud Director Tenant Portal, select an organization, and click **Go to vApps**.
4. In the **Virtual Applications** tab, in the bottom left of the vApp, select **Actions** > **Power** > **Start**.

The VMware Cloud Director Container Service Extension server is powered on.

- Create, manage, and upgrade Tanzu Kubernetes Grid clusters.
- Share VMware Cloud Director Container Service Extension catalogs with Non-administrative tenant users.

Update the VMware Cloud Director Container Service Extension Server

You can update the VMware Cloud Director Container Service Extension server in the **Server Details** section of the **CSE Management** tab in the Kubernetes Container Clusters UI plug-in.

Ensure you have an existing configured VMware Cloud Director Container Service Extension server. For more details, see Configure the VMware Cloud Director Container Service Extension Server.

1. In VMware Cloud Director UI, from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. In Kubernetes Container Clusters UI plug-in, select **CSE Management** > **Server Details** > **Update Server**.
3. In the **Update CSE Server** window, select one of the following specific workflows in each server upgrade process: Update Configuration, Patch Version Upgrade, Minor Version Upgrade.

   NOTE
   Once the VMware Cloud Director Container Service Extension server is fully setup and configured, if the service provider attempts to upgrade the VMware Cloud Director version to

10.4.2, the following alert appears in the **Getting Started** page in Kubernetes Container Clusters

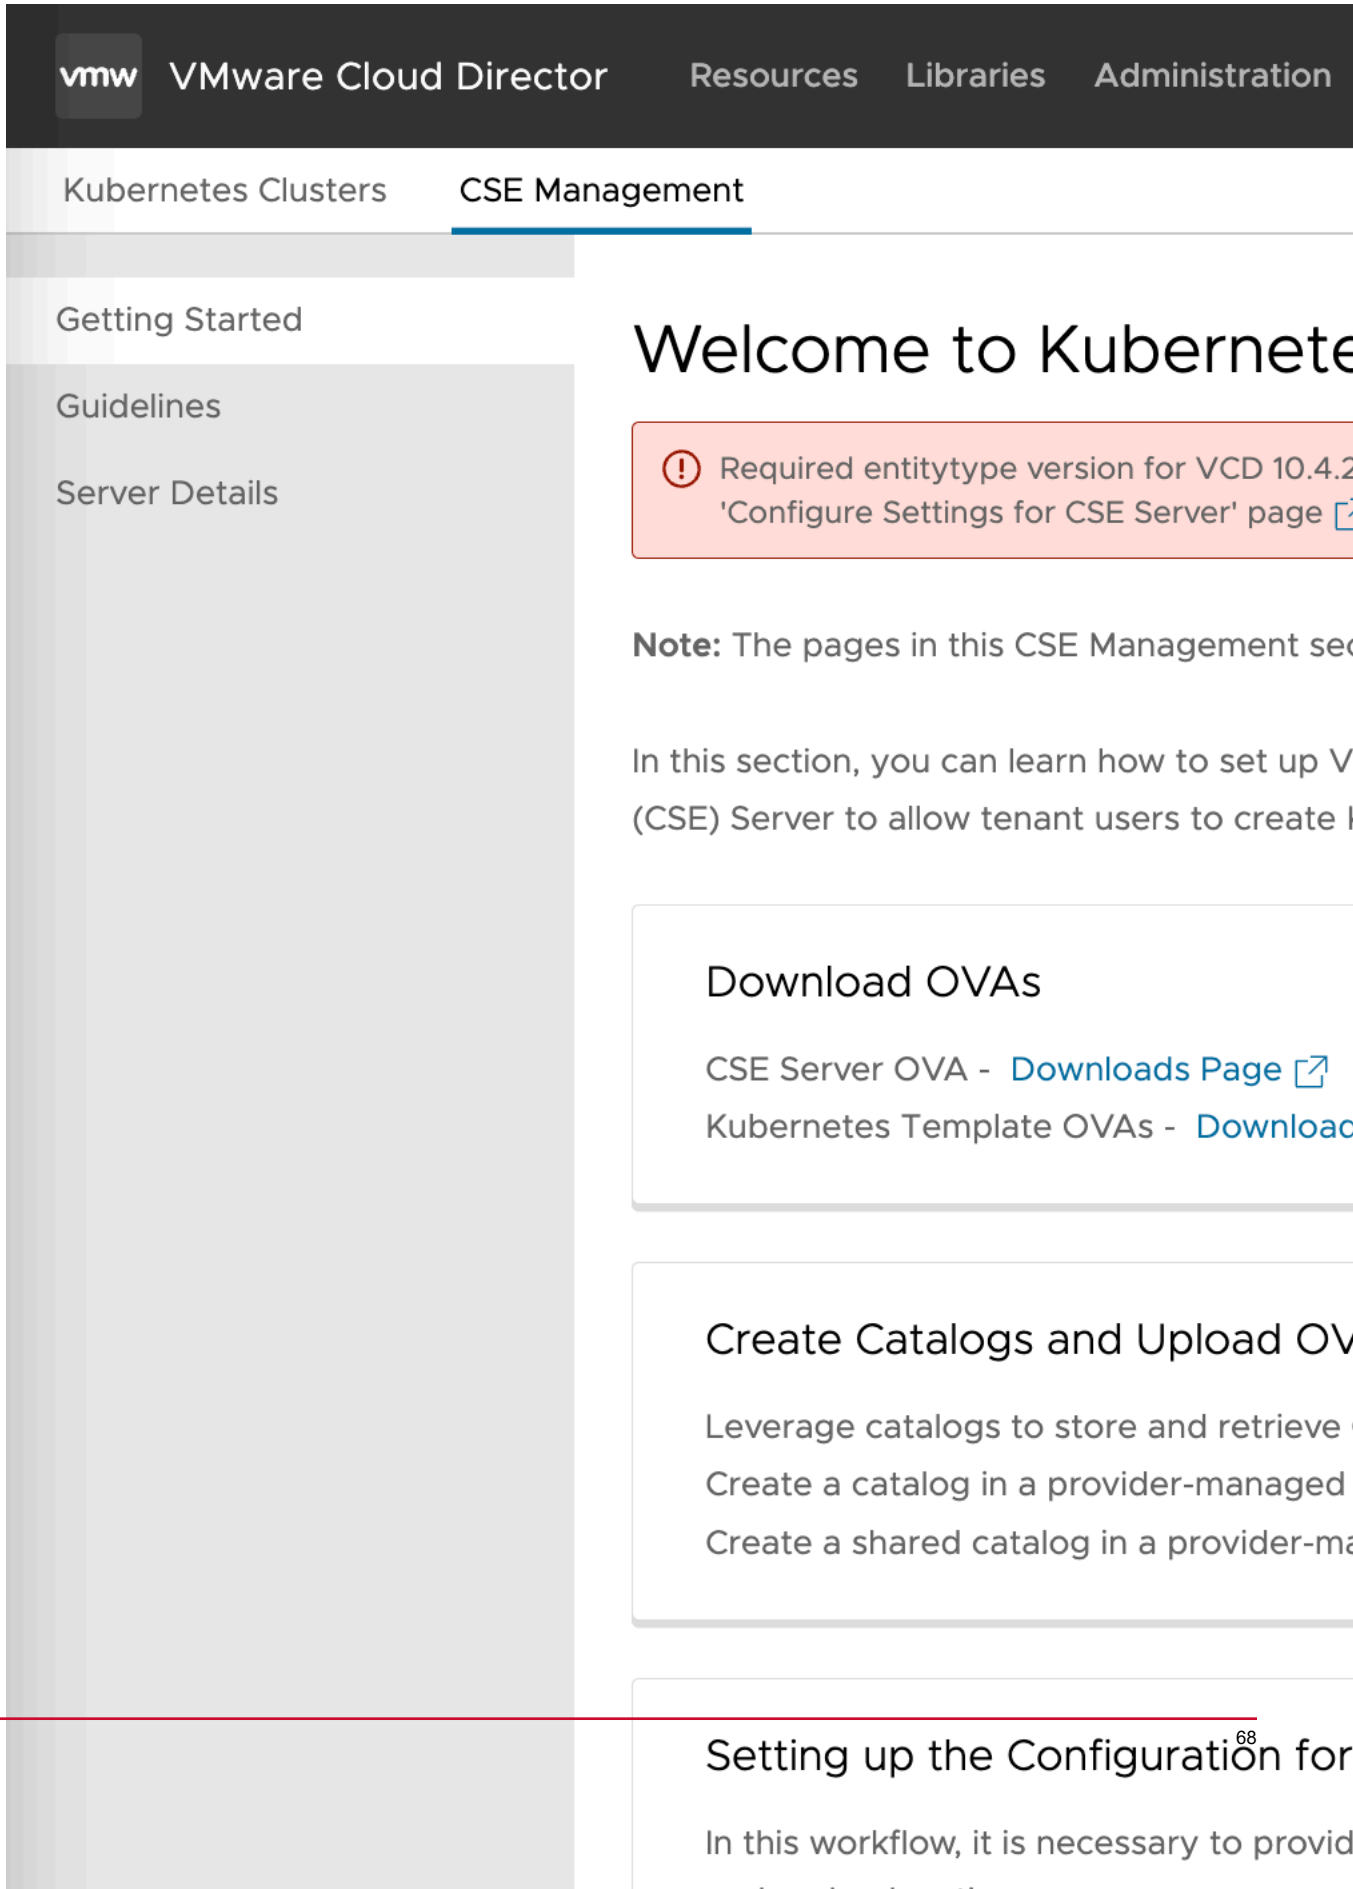

To rectify this error, complete the following steps:

1. In the **Setting up the Configuration for CSE Server** section, click Start.
2. In the **Configure Settings for CSE Server** window, in the **Set Up Prerequisites section**, click **Start** to register the required entitytype and other associated data.

After this error alert is rectified, you may proceed normally with the installation, or minor version upgrade workflow.

Update Server Configuration

In the Update Server Configuration process, you can add changes to the existing server configuration entity.

1. In the **Current CSE Server Components** section, edit the existing server configuration. For more information on the configuration parameters in this section, see Server Details.

   NOTE
   To restore the previous values of the form, click **Restore Previous**.

2. Click **Submit Changes**.

   NOTE
   These changes do not effect existing clusters. To apply these changes to existing clusters, it is necessary to manually update the cluster configuration.

3. Restart the existing VMware Cloud Director Container Service Extension Server vApp to apply the updated configuration.

   The server configuration has been updated and the new configuration will be applied to new clusters.

Minor Version Upgrade

To upgrade VMware Cloud Director Container Service Extension from one minor version to another, you can use the Minor Version Upgrade workflow in Kubernetes Container Clusters UI.

Take note of the Kubernetes component versions that were setup in VMware Cloud Director Container Service Extension configuration in the environment prior to upgrade. This is necessary if you want to retain the Kubernetes component

versions you used prior to the upgrade as this workflow automatically updates the Kubernetes components to the recommended versions.

To access the Minor Version Upgrade workflow, access Kubernetes Container Clusters UI as a service provider, click **CSE Management** > **Server Details** > **Update Server** > **Minor Version Upgrade**.

1. Manually stop and delete all existing VMware Cloud Director Container Service Extension Server vApps. For more information, refer to Power off a vApp and Delete a vApp.
2. Download the latest VMware Cloud Director Container Service Extension server OVA file. For more information, see Download OVA Files.
3. Start the **Set Up Prerequisites** automatic workflow.

Set up Prerequisite Step	Description
Register `Entitytypes`	`vcdkeconfig` 1.1.0 entitytype and capvcdcluster 1.3.0 or 1.1.0 is registered, depending on your VMware Cloud Director version. `capvcdcluster` 1.1.0 is registered if VMware Cloud Director version is lower than 10.4.2.
Update and publish**Kubernetes Clusters Rights Bundle**	Rights are added to this rights bundle, and then the rights bundle is published to all tenants. For a list of rights in this right bundle, see Kubernetes Clusters Rights Bundle.
Update and publish**Kubernetes Cluster Author** global role	Rights are added to this global role, and then the global role is published to all tenants. For a list of rights in this role, see Kubernetes Cluster Author Role.
Update Server Configuration	Automatically update Kubernetes component versions to the new recommended default values.

4. Optional: In the **Set Configuration Parameters** window, update the existing server configuration parameters, and click **Submit Changes**.

   NOTE
   The minor version upgrade workflow automatically updates Kubernetes components versions to the recommended versions in the VMware Cloud Director Container Service Extension configuration. If you want to continue using the component versions that you have been using before this minor version upgrade, you must provide those versions in the **Set Configuration Parameters** step.

   NOTE
   To restore the previous values of the form, click **Restore Previous**.

5. Optional: If you use a private registry to leverage an air-gapped environment, it is necessary to setup a local container registry using the `getting-started` script based upon your VMware Cloud Director Container Service Extension version before you can perform a minor version upgrade in VMware Cloud Director Container Service Extension. For more information, see Set up a Local Container Registry in an Air-gapped Environment.
6. Deploy a new vApp using the downloaded OVA. For more information, see Create a vApp from VMware Cloud Director Container Service Extension server OVA file.

Patch Version Upgrade

Learn how to upgrade the VMware Cloud Director Container Service Extension server to a higher patch version.

This workflow does not automatically update Kubernetes component versions to recommended versions. You can manually input component versions listed in the Compatibility section, or if the versions you currently use are higher than recommended versions, you can retain the higher versions.

A patch upgrade of VMware Cloud Director Container Service Extension works with a replacement model of VMware Cloud Director Container Service Extension vApp. It is not an in-place upgrade of the VMware Cloud Director Container Service Extension vApp. Additionally, you can also update the server configuration in this patch version upgrade workflow.

During this process, tenant users can continue to use the Kubernetes Container Clusters plug-in, and perform cluster operations. The cluster operations queue in VMware Cloud Director until the patch version upgrade is complete, and the new VMware Cloud Director Container Service Extension vApp then processes the operations.

1. Manually stop and delete the existing VMware Cloud Director Container Service Extension server vApp
2. Download a VMware Cloud Director Container Service Extension server OVA that has a higher patch version. For more information, see Download OVA Files.
3. Optional: In the **Set Configuration Parameters** window, update the existing server configuration parameters, and click **Submit Changes**.

   NOTE
   To restore the previous values of the form, click **Restore Previous**.
4. Optional: If you use a private registry to leverage an air-gapped environment, it is necessary to setup your local container registry using the `getting-started` script based upon your VMware Cloud Director Container Service Extension version before you can perform a patch version upgrade in VMware Cloud Director Container Service Extension. For more information, see Set up a Local Container Registry in an Air-gapped Environment.
5. Deploy a new vApp using the downloaded OVA. For more information, see Create a vApp from VMware Cloud Director Container Service Extension server OVA file.

Tanzu Kubernetes Grid Templates

VMware Cloud Director Container Service Extension uses Tanzu Kubernetes Grid templates as building blocks for deployment of Kubernetes clusters. This section details how to upload the Tanzu Kubernetes Grid OVA files, and how to share the them with non-administrative tenant users.

Download Tanzu Kubernetes Grid Templates

With VMware Cloud Director Container Service Extension you can deploy clusters from Tanzu Kubernetes Grid templates.

As a service provider, you must download Tanzu Kubernetes Grid templates from the download page and upload them into VMware Cloud Director so tenant users can use them to create clusters using VMware Cloud Director Container Service Extension. After the templates are uploaded to VMware Cloud Director, tenants can use the Kubernetes Container Clusters plug-in to deploy clusters from the Tanzu Kubernetes Grid templates. For more information, see Create Catalogs and Upload OVA Files.

Sharing Tanzu Kubernetes Grid Templates

Service providers must share the catalog that holds the Tanzu Kubernetes Grid template OVA files, with non-administrative tenant users. This action allows non-administrative tenant users to create, manage, and upgrade clusters.

To share the catalog that holds Tanzu Kubernetes Grid template OVA files with tenants as read-only, you must login to VMware Cloud Director tenant portal as a service provider. For sharing instructions, see Share a Catalog.

For more information on creating this catalog, see Create Catalogs and Upload OVA Files.

Using VMware Cloud Director Container Service Extension as a Service Provider

This guide provides information about how to use VMware Cloud Director™ Container Service Extension™ as a service provider. Service providers can manage tenant user rights, create and manage VMware Tanzu® Kubernetes Grid™ templates and clusters.

Intended Audience

This guide is intended for **Service Providers** to prepare tenant organization environments to use VMware Cloud Director Container Service Extension.

VMware Cloud Director Container Service Extension Requirements for Service Providers

To use VMware Cloud Director Container Service Extension 4.0 and later, ensure you are satisfying the following prerequisites.

Service Provider Requirements

- Ensure the VMware Cloud Director Container Service Extension server that resides in the solution organization can reach VMware Cloud Director load balancer endpoint.
- Ensure you are using the following network configuration.
  - Use the reference architecture to configure NSX and NSX Advanced Load Balancer correctly.
  - Deploy a test virtual service in a tenant organization to test the NSX and NSX Advanced Load Balancer configuration before you allow tenant users to begin cluster creation.
  - Ensure MTU (9000) values are correctly set on NSX, VMware ESX® VMkernel, adapters, and NSX Advanced Load Balancer.
  - Ensure MTU (9000) configuration is set correctly for the VMware Cloud Director Container Service Extension server to communicate to VMware Cloud Director load balancer endpoint. For more information, see https://knowledge.broadcom.com/external/article?legacyId=90850.
  - Ensure enough NSX Advanced Load Balancer licenses are available.
  - Ensure you are not using the 172.17.0.0/16 and 172.18.0.0/16 CIDR ranges or IP addresses from these ranges in the following network assets. These CIDR ranges are reserved by Docker and are used during the creation of bootstrap clusters.
    - Organization VDC network ranges where your Tanzu Kubernetes Grid clusters are deployed.
    - External IP allocations and ranges that are used by the Organization Edge Gateway and the associated Load Balancer.
    - Infrastructure networks where your DNS servers are connected.
    - The IP address, which the VMware Cloud Director public API endpoint URL resolves to.
- Ensure the VMware Cloud Director Container Service Extension server started successfully. Log into the VMware Cloud Director Container Service Extension server, and use one of the following commands to check the server status: `systemctl status cse.service` or `cse.log`.
- Ensure you input Github personal access token to avoid github API rate limit errors during cluster creation. Otherwise, cluster creation fails, particularly in concurrent attempts. For air-gapped environments, dot not input a Github personal access token.
- For a setup with multiple vCenter instances in VMware Cloud Director, ensure the Tanzu Kubernetes Grid OVA catalog syncs across the vCenter instances so that latency is not problematic during cluster operations.

- Ensure you manually update Custom roles that are cloned from the **CSE Admin Role** or **Kubernetes Cluster Author** role.

Organization Virtual Data Center Prerequisites for Kubernetes Cluster Deployment

This section details the prerequisites that are necessary for organization virtual data centers where tenant users create clusters.

Tenant organization virtual data centers require a routed, NSX backed organization virtual data center network. This network must have Internet connectivity to allow cluster VMs to download packages during configuration. It can be a routed organization VDC network with a source network address translation (SNAT) rule that allows VMs connected to the network access to outside world.

The edge gateway that powers the network must have the following settings:

- Free static IP addresses that are assigned to load balancers that front the Kubernetes cluster's control plane nodes.
- NSX Load Balancer activated, and Service Engine group assigned to it

User Roles and Rights

You can assign different roles, and right bundles to VMware Cloud Director Container Service Extension users. Each right and role allows users to perform different actions such as life cycle management of Tanzu Kubernetes Grid clusters, and perform administrative tasks in the Kubernetes Container Clusters UI.

The following table details the **Kubernetes Clusters Rights Bundle**, **Kubernetes Cluster Author** role, and **CSE Admin Role** that are created during the VMware Cloud Director Container Service Extension server configuration process.

Right Bundle or Role	Description
Kubernetes Clusters Rights Bundle	This right bundle comprises of the rights required for managing Tanzu Kubernetes Grid clusters. By default, this right bundle is automatically published to all tenants. Service providers have the ability to publish and unpublish this rights bundle to specific tenants afterwards. For more information, see Publish or Unpublish a Rights Bundle.
Kubernetes Cluster Author role	Assign this role to a user to manage Kubernetes clusters. For more information, see Assign Kubernetes Cluster Author Role to Tenant Users. In order for organization administrators to view all of the clusters in an organization, it is necessary to grant the user the **Administrator View: VMWARE:CAPVCDCLUSTER** right. Service providers must inform organization administrators that existing VMware Cloud Director Container Service Extension tenant users must be reassigned to the new **Kubernetes Cluster Author** role that is created in VMware Cloud Director Container Service Extension.
CSE Admin Role	You can create a user with **CSE Admin Role** in the system organization. The **CSE Admin Role** allows a user to perform administrative tasks in VMware Cloud Director Container Service Extension. You can use these user credentials as OVA deployment parameters when you start the VMware Cloud Director Container Service Extension server. For information, see Create a User with CSE Admin Role.

You can view the specific rights that are included in the **Kubernetes Clusters Rights Bundle**, **Kubernetes Cluster Author** role, and **CSE Admin Role** in the following sections.

Kubernetes Clusters Rights Bundle

The Kubernetes Clusters rights bundle comprises of the rights required for managing Tanzu Kubernetes Grid clusters. By default, this right bundle is automatically published to all tenant users of VMware Cloud Director Container Service Extension.

Table 9: Rights included in Kubernetes Clusters Rights Bundle

Right
Allow Access to All Organization VDCs
Manage user's own API token
Manage Certificates Library
View Certificates Library
Administrator View
Create a Disk
Edit Disk Properties
View Disk Properties
Create a Shared Disk
Preserve All ExtraConfig Elements During OVF Import and Export
View Shared Catalogs from Other Organizations
View Gateway
NAT View Only
NAT Configure
Load Balancer View Only
Load Balancer Configure
View: VMWARE:VCDKECONFIG
View: VMWARE:CAPVCDCLUSTER
Edit VMWARE:CAPVCDCLUSTER
Full Control: VMWARE:CAPVCDCLUSTER
Administrator View: VMWARE:CAPVCDCLUSTER
Administrator Full Control: VMWARE:CAPVCDCLUSTER

The following IP Spaces rights are optional, and are only necessary when you want to leverage Gateways using IP Spaces. Service providers must manually add these rights to the **Kubernetes Clusters Rights Bundle** as they are not automatically added. For instructions, see View and Edit a Rights Bundle Using VMware Cloud Director.

Table 10: IP Spaces Rights

Right
View IP Spaces
Manage IP Spaces
Allocate IP Spaces

The following conditional rights are added only if they already exist in the system, as they are relevant to Kubernetes cluster management:

Table 11: Conditional Rights

Right
View: Tanzu Kubernetes Guest Cluster
Edit Tanzu Kubernetes Guest Cluster
Full Control: Tanzu Kubernetes Guest Cluster
Administrator View: Tanzu Kubernetes Guest Cluster
Administrator Full Control: Tanzu Kubernetes Guest Cluster
View: CSE:NATIVECLUSTER
Edit CSE:NATIVECLUSTER
Full Control: CSE:NATIVECLUSTER
Administrator View: CSE:NATIVECLUSTER
Administrator Full Control: CSE:NATIVECLUSTER

Kubernetes Cluster Author Role

A tenant user with the **Kubernetes Cluster Author** role can view, create, and manage Kubernetes clusters. This role simplifies tenant role setup by assigning all the required rights in one role in the VMware Cloud Director UI.

Table 12: Rights included in the Kubernetes Cluster Author Role

Right
Allow Access to All Organization VDCs
View Organization Administrative Details
View vApp ACL
Manage user's own API token
View Certificates Library
View Compute Policies for an Organization VDC
View Disk IOPS
View Disk Encryption Status
View Disk Properties
Create a Disk

Right
Delete a Disk
Edit Disk Properties
Create a Shared Disk
Edit VM-VM Affinity Rule
View Encryption Status of VMs and VM's disks
View VM metrics
Preserve All ExtraConfig Elements During OVF Import and Export
Copy a vApp
Create / Reconfigure a vApp
Delete a vApp
Download a vApp
Edit vApp Properties
Edit VM Compute Policy
Edit VM CPU
Edit VM Hard Disk
Edit VM Memory
Edit VM Network
Edit VM Properties
Manage VM Password Settings
Start / Stop / Suspend / Reset a vApp
Share a vApp
Create / Revert / Remove a Snapshot
Upload a vApp
Access to VM Console
Edit / View VM Boot Options
Allow metadata mapping domain to vCenter
View Tenant Portal Plugin Information
View Shared Catalogs from Other Organizations
View Private and Shared Catalogs within Current Organization
Add a vApp from My Cloud
View vApp Templates / Media
Copy / Move a vApp Template / Media
Edit vApp Template / Media Properties
Add to My Cloud
View Gateway
NAT View Only
NAT Configure

Right
Load Balancer View Only
Load Balancer Configure
View Properties
View: VMWARE:CAPVCDCLUSTER
Edit VMWARE:CAPVCDCLUSTER
Full Control: VMWARE:CAPVCDCLUSTER
View: VMWARE:VCDKECONFIG

The following IP Spaces rights are optional, and are only necessary when you want to leverage Gateways using IP Spaces. Service providers must manually add these rights to the **Kubernetes Cluster Author** role as they are not automatically added. For instructions, see View and Edit a Global Tenant Role Using VMware Cloud Director.

Table 13: IP Spaces Rights

Right
View IP Spaces
Manage IP Spaces
Allocate IP Spaces

The following conditional rights are added only if they already exist in the system, as they are relevant to Kubernetes cluster management:

Table 14: Conditional Rights

Right
View: Tanzu Kubernetes Guest Cluster
Edit Tanzu Kubernetes Guest Cluster
Full Control: Tanzu Kubernetes Guest Cluster
View: CSE:NATIVECLUSTER
Edit CSE:NATIVECLUSTER
Full Control: CSE:NATIVECLUSTER

CSE Admin Role

The **CSE Admin Role** contains rights that allow a service provider or organization administrator to perform administrative tasks in VMware Cloud Director Container Service Extension.

To learn how to create a user with the **CSE Admin Role**, see Create a User with the CSE Admin Role.

Table 15: Rights included in the CSE Admin Role

Right
Manage user's own API token
View: VMWARE:VCDKECONFIG
Edit VMWARE:VCDKECONFIG
Full Control: VMWARE:VCDKECONFIG
Administrator View: VMWARE:VCDKECONFIG
Administrator Full Control: VMWARE:VCDKECONFIG
View: VMWARE:CAPVCDCLUSTER
Edit VMWARE:CAPVCDCLUSTER
Full Control: VMWARE:CAPVCDCLUSTER
Administrator View: VMWARE:CAPVCDCLUSTER
Administrator Full Control: VMWARE:CAPVCDCLUSTER

Check the VMware Cloud Director Container Service Extension Server Status

Follow these steps to check the VMware Cloud Director Container Service Extension service status:

Ensure you configure the VMware Cloud Director Container Service Extension server. For more information, see Configure the VMware Cloud Director Container Service Extension Server.

1. Log into the virtual machine that runs the VMware Cloud Director Container Service Extension OVA file.
2. To check the VMware Cloud Director Container Service Extension server status, run the following command:

   ```
   systemctl status cse
   ```

   The status of the VMware Cloud Director Container Service Extension server appears.

Working with Kubernetes Clusters

Learn how to create, configure, delete, and upgrade Kubernetes clusters by using VMware Cloud Director Container Service Extension. The primary tool for these operations is Kubernetes Container Clusters plug-in that accompanies VMware Cloud Director.

Create a Tanzu Kubernetes Grid Cluster

Starting with VMware Cloud Director 10.3.1, you can create Tanzu Kubernetes Grid clusters by using the Kubernetes Container Clusters UI plug-in.

- Ensure that you have set up, and powered on the VMware Cloud Director Container Service Extension server. For more information, see Configure the VMware Cloud Director Container Service Extension Server.
- Ensure that you are satisfying the VMware Cloud Director Container Service Extension requirements for service providers. See VMware Cloud Director Container Service Extension Requirements for Service Providers.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters** > **New**.
2. Select the **VMware Tanzu Kubernetes Grid** runtime option, and click **Next**.
3. Enter a name, select a Kubernetes template from the list, and click **Next**.
4. In the **VDC & Network** window, select the organization VDC to which you want to deploy a Tanzu Kubernetes Grid cluster, select a VDC network for the cluster, and click **Next**.
5. In **Control Plane** window, select the number of nodes, disk size, optionally select a sizing policy, a placement policy, a storage profile, and click **Next**.

   NOTE
   The number of nodes input allows for clusters to have multiple control plane nodes.
6. In **Worker Pools** window, enter a name, number of nodes, disk size, optionally select a sizing policy, a placement policy, a storage profile, and click **Next**. For more information on worker node pools, see Working with Worker Node Pools.

   NOTE

   - To configure vGPU settings, select the **Activate GPU** toggle and select a vGPU policy. For more information on vGPU configuration, see Configuring vGPU on Tanzu Kubernetes Grid Clusters to allow AI and ML Workloads.
   - When you create clusters with vGPU functionality, it is recommended to increase the disk size to between 40-50 GB as vGPU libraries occupy a large amount storage space.
   - You can select a sizing policy in this workflow or separately in VMware Cloud Director Container Service Extension server configuration. When you select a sizing policy in conjunction with a vGPU policy that contains VM sizing, the sizing information in the vGPU policy takes precedence over the selected sizing policy. It is recommended to include sizing in your vGPU policy, and only specify a vGPU policy when you leave the **Sizing Policy** field empty.

7. Optional: To create additional worker node pools, click **Add New Worker Pool**, and configure worker node pool settings.
8. Click **Next**.
9. In the **Kubernetes Storage** window, activate the **Create Default Storage Class** toggle, select a storage profile and enter a storage class name.
10. Optional: Configure **Reclaim Policy** and **Filesystem** settings.
11. In the **Kubernetes Network** window, specify a range of IP addresses for Kubernetes services and a range for Kubernetes pods, and click **Next**.

    Classless Inter-Domain Routing (CIDR) is a method for IP routing and IP address allocation.

Option	Description
Pods CIDR	Specifies a range of IP addresses to use for Kubernetes pods. The default value is 100.96.0.0/11. The pods subnet size must be equal to or larger than /24. You can enter one IP range.
Services CIDR	Specifies a range of IP addresses to use for Kubernetes services. The default value is 100.64.0.0/13. You can enter one IP range.
Control Plane IP	Tenant users can specify their own IP address as the control plane endpoint. They can use an external IP address from the gateway or an internal IP address from a subnet that is different from the routed IP range. If they do not specify an IP address as the control plane endpoint, VMware Cloud Director Container Service Extension server selects one of the unused IP addresses from the associated tenant gateway.
Virtual IP Subnet	Tenant users can specify a subnet CIDR from which one unused IP address is assigned as Control Plane Endpoint. The subnet must represent a set of addresses that are present in the gateway. The same CIDR is also propagated as the subnet CIDR for the ingress services on the cluster.

You can use the following IP addresses as the Control Plane IP:

IP Type	Description
External IP addresses	Any of the IP addresses in the external gateway that connect to the OVDC network.
Internal IP addresses	Any private IP address that is internal to the tenant, with the following exceptions: • IP addresses in the LB network service definition, usually 192.168.255.1/24. • IP addresses that are in the organization VDC IP subnet. • IP address that is in use.

NOTE
When an IP address does not have the above characteristics, the following behavior occurs:

- If the IP address is already in use, and VMware Cloud Director detects the usage, an error appears in the logs during LB creation.
- If the IP address is already in use, and VMware Cloud Director does not detect the usage, the behavior is undefined.

12. In the **Debug Settings** window, activate or deactivate the **Auto Repair on Errors** toggle, and the **Node Health Check** toggle.

Toggle	Description
Auto Repair on Errors	This toggle applies to failures that occur during the cluster creation process. If you activate this toggle, the VMware Cloud Director Container Service Extension server attempts to recreate the clusters that are in an error state during the cluster creation process. If you deactivate this toggle, the VMware Cloud Director Container Service Extension server leaves the cluster in an error state for manual troubleshooting. **NOTE** This toggle is deactivated by default in VMware Cloud Director Container Service Extension 4.1.and newer versions. Service providers must advise tenant users of this as it is a behavioral change from VMware Cloud Director Container Service Extension 4.0.
Node Health Check	In contrast to Auto Repair on Errors when the remediation process is only applicable during cluster creation, the remediation process in Node Health Check begins after the cluster reaches an available state. If any of the nodes become unhealthy during the life time of the cluster, Node Health Check detects and remediates them. For more information, see Node Health Check Configuration. **NOTE** This toggle is deactivated by default in VMware Cloud Director Container Service Extension 4.2.

13. Enter an SSH public key.
14. Click **Next**.
15. Review the cluster settings and click **Finish**.

    NOTE
    In the **Review** window, a warning appears to advise you that the cluster contains an API token of the owner, and not to share the kubeconfig or the cluster directly with others. Instead, create the cluster as a tenant user of an organization.

Review Cluster Status

When you create a Tanzu Kubernetes Grid cluster in VMware Cloud Director Container Service Extension, the following status appear:

Table 16: Cluster Status

Cluster Status	Description
Pending	The cluster request has not yet been processed by the VMware Cloud Director Container Service Extension server.
Creating	The cluster is currently being processed by the VMware Cloud Director Container Service Extension server.
Available	The cluster is ready for users to operate on and host workloads.

Cluster Status	Description
Deleting	The cluster is being deleted
Error	The cluster is in an error state. **NOTE** If you want to manually debug a cluster, deactivate **Auto Repair on Errors** mode.

View Tanzu Kubernetes Grid Cluster Information

This section details how to view the configuration information of a Tanzu Kubernetes Grid cluster in the Kubernetes Container Clusters UI plug-in.

You can view the following sections in the cluster information page:

Table 17: Cluster Information Page Sections

Tabs	Description
Overview	This tab details the overall configuration of the Tanzu Kubernetes Grid cluster: • **Info**: Basic cluster information such as cluster name, status, and Kubernetes version. • **Kubernetes Resources**: CAPVCD version, Cluster Resource Set Bindings, CPI, CSI • **vApp Details**: Virtual Data Center, Network, Owner, Cluster ID
Node Pools	This tab details the node pools that exist in the cluster. For more information on node pools, see Working with Worker Node Pools.
Kubernetes Storage	This tab details Kubernetes default storage class configuration and persistent volumes. For more information, see Configure a Default Storage Class and Working with Stateful Deployments.
Events	This tab details each event that occurs in the Tanzu Kubernetes Grid cluster after creation. Click each event to view event details, such as the event name, event type, event time, and resource name.

1. To view Tanzu Kubernetes Grid cluster information, in VMware Cloud Director UI, from the top navigation bar, select **More>Kubernetes Container Clusters**.
2. In the **Kubernetes Clusters** tab, in the datagrid, click on the name of the cluster you want to view.

The Tanzu Kubernetes Grid cluster information page appears, and you can navigate through each tab to view specific cluster information.

Working with Worker Node Pools

In VMware Cloud Director Container Service Extension, you can create, resize and delete worker node pools in the Kubernetes Container Clusters UI plug-in in VMware Cloud Director.

Worker node pools are groups of worker nodes in a cluster that share the same configuration on which your workloads can run on. By configuring worker node pools, clusters can have several different types of worker nodes to perform separate tasks in the one cluster. It is necessary to have at least one worker node pool with one worker node for a Tanzu Kubernetes Grid cluster.

Create a Worker Node Pool

In VMware Cloud Director Container Service Extension, you can create multiple worker node pools in a Tanzu Kubernetes Grid cluster. Follow these steps to create a worker node pool or add additional worker node pools to an existing cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. From the list of clusters, click the name of the cluster to which you want to add a worker node pool.
3. In the cluster information window, click the **Node Pools** tab and click **Create Worker Node Pools**.
4. In the **Create New Worker Node Pool** window, enter a worker node pool name, number of nodes and disk size, and optionally select a sizing policy, placement policy and storage policy.

   NOTE

   - To configure vGPU settings, select the **Activate GPU** toggle. For more information on vGPU configuration, see Configuring vGPU on Tanzu Kubernetes Grid Clusters to allow AI and ML Workloads.
   - When you create clusters with vGPU functionality, it is recommended to increase the disk size to between 40-50 GB as vGPU libraries occupy a large amount storage space.
   - You can select a sizing policy in this workflow or seperately in VMware Cloud Director Container Service Extension server configuration. When you select a sizing policy in this workflow, it takes precedence in the Tanzu Kubernetes Grid cluster configuration.
5. Optional: Click **Create New Worker Pool** to create additional worker node pools.
6. Click **Create**.

Resize a Node Pool

In this section, you can learn how to resize an existing node pool by adding or reducing the number of worker nodes in a Tanzu Kubernetes Grid cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. From the list of clusters, click the name of the cluster to resize.
3. In the cluster information page, click the **Node Pools** tab.
4. On the left of the node pool you want to resize, click the ellipsis, and select **Resize**.
5. In the **Resize Node Pool** window, configure the number of nodes, and click **Submit**.

Working with Stateful Deployments

This section details how to work with stateful deployments using VMware Cloud Director Container Service Extension.

You can deploy stateful applications on VMware Cloud Director Container Service Extension 4.x provisioned Tanzu Kubernetes Grid clusters. The Tanzu Kubernetes Grid clusters include Container Storage Interface preinstalled, that activates both static and dynamic persistent volumes. For more information, see Dynamic Persistent Volumes. For more information on Container Storage Interface, refer to Container Storage Interface (CSI) driver for VMware Cloud Director Named Independent Disk.

Configure a Default Storage Class

Starting with VMware Cloud Director Container Service Extension 3.1.3, you can optionally configure a default storage class when you create a Tanzu Kubernetes Grid cluster, and this storage class is used by default for the creation of any persistent volumes. This feature automates steps for the tenant users to manage the Tanzu Kubernetes Grid cluster for their developers in their organization.

In the Kubernetes Container Clusters UI plug-in in VMware Cloud Director, the **Create Default Storage Class** toggle is activated by default. You can deactivate the toggle to opt out of the function.

You can configure the following fields in a default storage class:

Table 18: Default Storage Class Configuration Fields

Configuration field	Description
VMware Cloud Director Storage Profile Name	Select one of the available VMware Cloud Director storage profiles.
Storage Class Name	The name of the default Kubernetes storage class. This field can be any user-specified name with the following constraints, based on Kubernetes requirements: • Contain a maximum of 63 characters • Contain only lowercase alphanumeric characters or hyphens • Start with an alphabetic character • End with an alphanumeric character
Reclaim Policy	• **Delete** policy: This policy is set by default. This policy deletes the PersistentVolume object when the PersistentVolumeClaim is deleted. • **Retain** policy: This policy does not delete the volume when the PersistentVolumeClaim is deleted, and the volume can be reclaimed manually.
Filesystem	• **xfs** • **ext4**: This is the default filesystem used for the storage class.

For instructions on how to configure a default storage class in the Kubernetes Container Clusters UI plug-in in VMware Cloud Director, see Create a Tanzu Kubernetes Grid Cluster.

Configuring vGPU on Tanzu Kubernetes Grid Clusters to allow AI and ML Workloads

You can deploy artificial intelligence (AI) and machine learning (ML) workloads on clusters provisioned by the Tanzu Kubernetes Grid. The deployment of artificial intelligence and machine learning workloads requires some initial setup by service providers, and some configuration by organization administrators and tenant users in the cluster creation workflow.

To prepare VMware Cloud Director environment to provision clusters that can handle artificial intelligence and machine learning workloads, service providers must create a vGPU policy and add a vGPU policy to an organization VDC. For instructions on how to perform these tasks, refer to Creating and Managing vGPU Policies. Once service providers perform these steps, tenant users can deploy artificial intelligence and machine learning workloads to their Tanzu Kubernetes Grid clusters.

To create Tanzu Kubernetes Grid clusters with vGPU functionality, see Create a Tanzu Kubernetes Grid Cluster. If you are using Tanzu Kubernetes Grid 2.1 and above that are interoperable with VMware Cloud Director Container Service Extension, the following sections are not applicable and you can proceed to the cluster creation workflow.

> **NOTE**
> The following sections are applicable to Tanzu Kubernetes Grid 1.6.1 only, that is no longer supported by VMware. To avail of the vGPU functionality, use Tanzu Kubernetes Grid versions 2.1 and above that are interoperable with VMware Cloud Director Container Service Extension.

BIOS Firmware Limitations

VMware Cloud Director Container Service ExtensionTanzu Kubernetes Grid templates are built with BIOS firmware, and it is not possible to change this firmware configuration. The BAR1 memory on this firmware cannot exceed 256 MB. NVIDIA

Grid cards with more than 256MB of BAR1 memory require EFI firmware. For more information on firmware limitations, refer to VMware vSphere: NVIDIA Virtual GPU Software Documentation.

Create a Custom Image with EFI Firmware

To overcome the BIOS firmware limitations that exist on Tanzu Kubernetes Grid templates, you can create a custom image with EFI firmware in vSphere. For instructions, refer to **Linux Custom Machine Images** sections in the archived Tanzu Kubernetes Grid 1.6 documentation. To access the archived documentation, see VMware Tanzu Kubernetes Grid Documentation**> Unsupported Releases.**

To create Linux custom machine images with Tanzu Kubernetes Grid 1.6 successfully on a GPU template, you also have to include the following inputs when you build the custom image:

<table>
<tr><th>Inputs</th><th>Description</th></tr>
<tr><td><code>customizations.json</code></td><td>To build an image for a vGPU-enabled cluster for vSphere, create a file named <code>customizations.json</code> , and add the following:
<pre>{
"vmx_version": "17"
}</pre></td></tr>
<tr><td><code>metadata.json</code></td><td><code>VERSION</code> must identically match an established version of a Tanzu Kubernetes Grid template, as the Kubernetes Container Clusters UI plug-in does not recognize the OVA file if the version number differs to that of the template.
The following example outlines the recommended file naming convention:
<table><tr><th>Template and Version</th><th>Metadata</th></tr><tr><td>Kubernetes template for TKG 1.6</td><td><code>ubuntu-2004-kube-v1.23.10+vmware.1-tkg.2-b53d41690f8742e7388f2c553fd9a181.ova</code></td></tr><tr><td>Version</td><td><code>v1.23.10+vmware.1-tkg.2-b53d41690f8742e7388f2c553fd9a181</code></td></tr></table></td></tr>
<tr><td><code>build-node-ova-vsphere-ubuntu-2004-efi</code></td><td>Use this command to run the image builder for vGPU-enabled clusters. This command specifies to build the custom image with EFI firmware.</td></tr>
</table>

Service providers must set up a new catalog in VMware Cloud Director for vGPU templates, and upload the templates to this catalog. When a user wants to create a vGPU-enabled cluster, they can select this template in the cluster creation process, and it leverages the vGPUs in that cluster. For more information, see Create Catalogs and Upload OVA Files.

Working with Ingress Services on Tanzu Kubernetes Grid Clusters

In this section, you can learn about Ingress-based TCP/HTTP/HTTPS services you can deploy on Tanzu Kubernetes Grid clusters.

Ingress services expose TCP/HTTP/HTTPS routes from outside the cluster to services within the cluster. Traffic routing is controlled by rules defined on the Ingress resource. For more information on Ingress, refer to the Kubernetes website.

Tanzu Kubernetes Grid clusters has a Kubernetes Cloud Provider for VMware Cloud Director preinstalled. This feature allows you to deploy Ingress services for modern applications. In this process, it might be necessary to upload the associated certificates with a particular naming convention. For more information, see Creation of a LoadBalancer using a Third-Party Ingress.

Upgrade a Tanzu Kubernetes Cluster

This section details how to upgrade Kubernetes versions in a Tanzu Kubernetes cluster in VMware Cloud Director Container Service Extension.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. From the cluster list, select the cluster you want to upgrade.

   NOTE
   Once you select a cluster, if an upgraded Kubernetes version is available, the newer version appears in the **Upgrade** section of the cluster information. If there are no upgrades available, the **Upgrade** tab in the operations menu deactivates.

   NOTE
   The following warning alert appears for Tanzu Kubernetes Grid clusters during cluster upgrade workflow. If this warning appears, and the current versions of Kubernetes components in the cluster do not match the available versions for upgrades, follow the instructions in the Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension

Clusters workflow. Do not continue with the cluster upgrade workflow you are currently in.

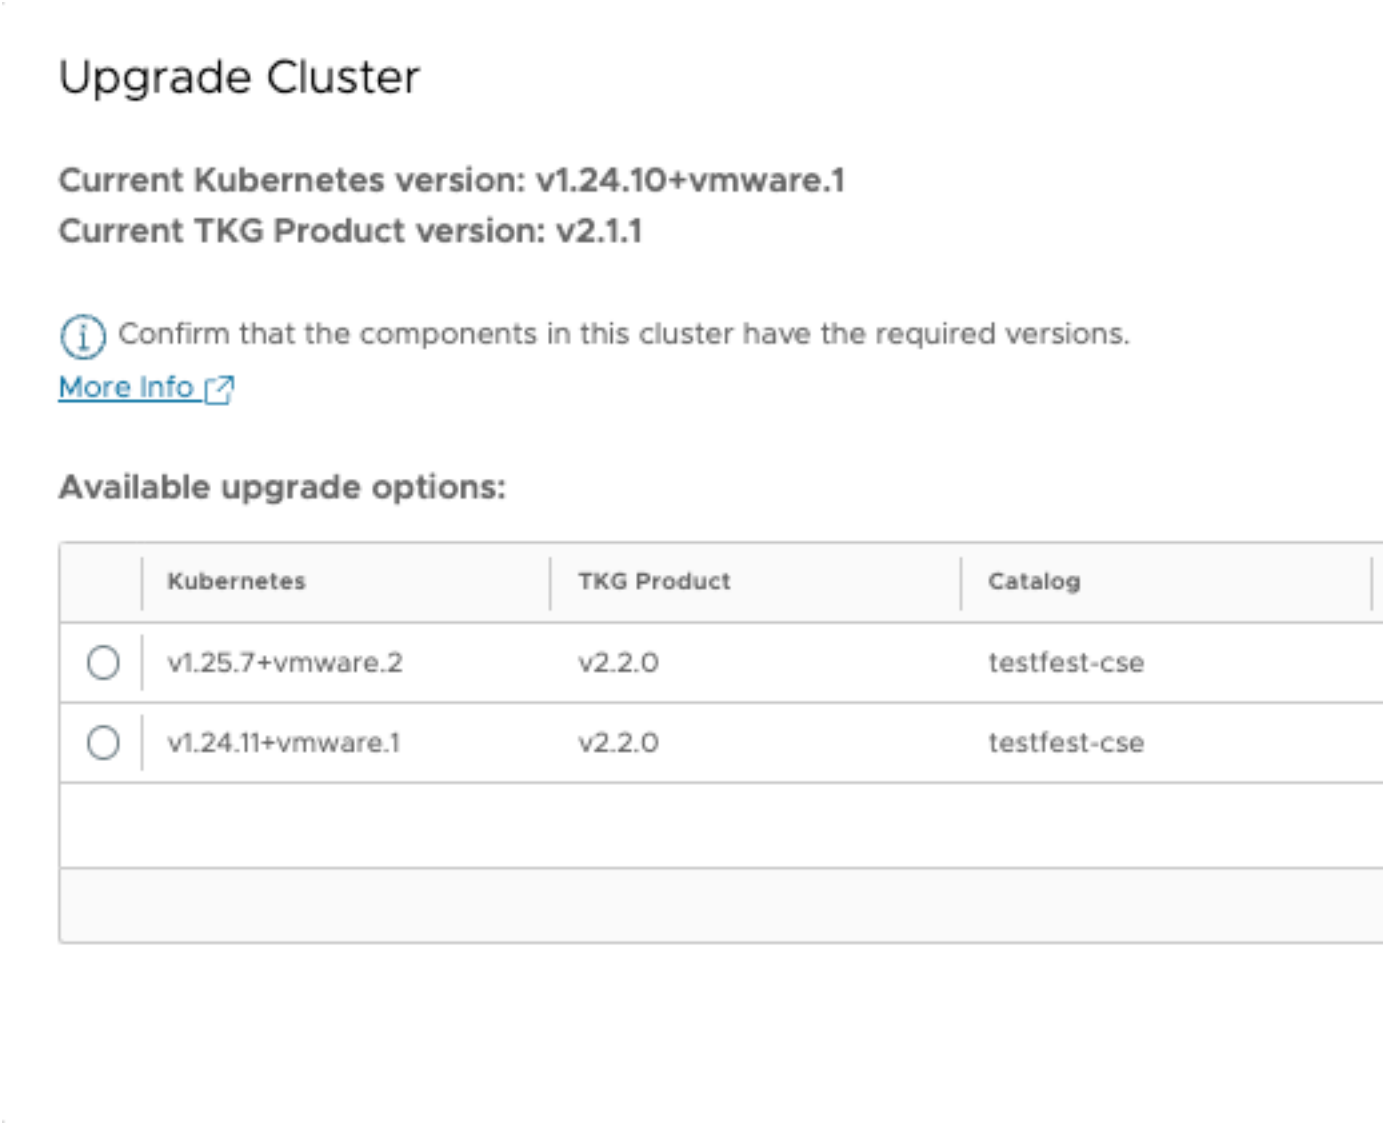

3. From the operations menu, click **Upgrade**.

   NOTE
   Once the upgrade is issued, you can see the upgrade spinner in the cluster list page for the associated cluster. The spinner is present until the upgrade is complete.

Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters

For clusters that were created using older versions of VMware Cloud Director Container Service Extension, it is necessary to perform a one time script upgrade. This allows the clusters to be compatible with the VMware Cloud Director Container Service Extension you want to use.

Before you perform this task, ensure the following prerequisites are in place:

- The machine must use Ubuntu GNU/Linux 20.04 operating system.
- Ensure kubectl is installed. For more information, see Install and Set Up kubectl on Linux.
- Ensure you install and operate Docker. For more information, see Install Docker Engine.
- Ensure the kubeconfig of the cluster is present on the machine at an accessible path.

By default, clusters created in older versions of VMware Cloud Director Container Service Extension, operate on the following older versions of the Kubernetes components. It is necessary to upgrade to newer versions outlined below.

Kubernetes Components	Existing Version	Upgrade Version for 4.2	Upgrade Version for 4.2.1	Upgrade Version for 4.2.2	Upgrade Version for 4.2.3
Kubernetes Cloud Provider for VMware Cloud Director	1.5.0, 1.4.1, or older versions	1.5.0	1.6.0	1.6.0	1.6.1
Kubernetes Container Storage Interface driver for VMware Cloud Director	1.5.0, 1.4.1, or older versions	1.5.0	1.6.0	1.6.0	1.6.0
Kubernetes Cluster API Provider for VMware Cloud Director	1.2.0, 1.1.1, or older versions	1.2.0	1.3.0	1.3.0	1.3.2
RDE-Projector	0.7.0, 0.6.1, or older versions	0.7.0	0.7.0	0.7.1	0.7.1

1. Use the following command to set `$HOME` directory:

   ```
   export $HOME=<directory of choice>
   ```

   NOTE
   To confirm `$HOME` directory is set correctly, use the following command to print the `$HOME` directory. It should not be empty.

   ```
   echo $HOME
   ```

2. Use the following command to create a folder structure for mounting and storing the cluster upgrade script content:

   ```
   mkdir -p $HOME/cluster-upgrade-script
   ```

3. Navigate to the folder created, and use the following command to pull the cluster-upgrade-script image from http://projects.packages.broadcom.com.

   - For 4.2:

     ```
     cd $HOME/cluster-upgrade-script
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.2
     ```

   - For 4.2.1:

     ```
     cd $HOME/cluster-upgrade-script
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.3
     ```

   - For 4.2.2:

     ```
     cd $HOME/cluster-upgrade-script
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.4
     ```

   - For 4.2.3:

     ```
     cd $HOME/cluster-upgrade-script
     ```

```
docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.5
```

4. After you pull the image, use docker to extract the image contents to the folder by creating a container and extracting it.
   - For 4.2:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.2
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
   - For 4.2.1:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.3
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
   - For 4.2.2:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.4
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
   - For 4.2.3:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.5
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
5. After the contents are extracted to the folder, it is safe to remove the temporary container:
   ```
   docker container rm "temp_container"
   ```
6. After all the content is extracted to the main folder, open the directory, and update the permissions for the script:
   ```
   cd $HOME/cluster-upgrade-script
   chmod +x upgrade_cluster_components.sh
   ```
7. Run the following command in the main folder:
   ```
   ./upgrade_cluster_components.sh kubeconfig-absolute-file-path [image registry]
   ```

```
Example: ./upgrade_cluster_components.sh $HOME/kubeconfig-cluster.txt projects.packages.broadcom.com
```

`image_registry` is an optional parameter, which defines the registry the script should pull images from and expects images such as CSI/CPI CRS, CAPVCD manifests, clusterctl, etc. to be hosted there.

When you are not using private/local/airgap registry, use `projects.packages.broadcom.com` as the `image_registry` . Alternatively if the parameter is empty, it defaults to `image_registry: projects.packages.broadcom.com.`

For private/local registry, use your registry link for the parameter. For example, `my-private.registry.com` . Ensure that the virtual machine that you run the scripts trust the registry, or it can run into errors such as `x509 certificate signed by unknown authority.`

NOTE

- During script execution, all image artifacts are downloaded to `$HOME/cluster-upgrade-packages/` .
- If there are any errors during the upgrade from the script, it is safe to remove this folder, and run the script again to create this folder.
- An additional folder is created for clusterctl at `$HOME/.cluster-api` . It is safe to delete this folder as re-attempting to run the script creates this folder.

After the upgrade cluster script has ran successfully, the cluster Kuberetes component versions are updated. You can view the updated Kubernetes component versions in the Kubernetes Container Clusters UI.

Resize a Tanzu Kubernetes Grid Cluster

This section details how to resize a Tanzu Kubernetes Grid cluster in VMware Cloud Director Container Service Extension.

To resize a Tanzu Kubernetes Grid cluster, it is necessary to resize the node pools within a cluster in Kubernetes Container Clusters UI plug-in.

Resize a Node Pool

In this section, you can learn how to resize an existing node pool by adding or reducing the number of worker nodes in a Tanzu Kubernetes Grid cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. From the list of clusters, click the name of the cluster to resize.
3. In the cluster information page, click the **Node Pools** tab.
4. On the left of the node pool you want to resize, click the ellipsis, and select **Resize**.
5. In the **Resize Node Pool** window, configure the number of nodes, and click **Submit**.

Delete a Kubernetes Cluster

This section details how to delete a Kubernetes cluster in VMware Cloud Director Container Service Extension.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Select a cluster, and click **Delete**.

The Kubernetes cluster is deleted.

Force Delete a Kubernetes Cluster

In VMware Cloud Director Container Service Extension, you can force delete Kubernetes clusters, and their associated resources that are not fully complete and that are in an unremovable state.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Select a cluster, and click **Delete**.
3. In the **Delete Cluster** window, click **Force Delete**, and **Delete**.

Any incomplete or previously unremovable Kubernetes clusters are deleted.

FAQs

This section details frequently asked questions about VMware Cloud Director Container Service Extension.

As a previous Kubernetes Container Clusters UI user, how does this new release affect me?

With this Kubernetes Container Clusters UI plug-in, users can create, manage, and upgrade Tanzu Kubernetes Grid clusters. The clusters are created with Cluster APIs of VMware Cloud Director, behind the scenes. All of your previously created clusters of types such as native, vSphere with Tanzu, and Tanzu Kubernetes Grid legacy are still available.

What is Tanzu Kubernetes Grid Clusters Legacy?

All Tanzu Kubernetes Grid clusters created using VMware Cloud Director Container Service Extension server 3.1.x are now referred to as Tanzu Kubernetes Grid Legacy clusters in 4.0 and newer versions.

As a first time user of Kubernetes Container Clusters 4.x, what do I need to know?

To allow tenant users to create, manage, and upgrade Tanzu Kubernetes Grid clusters, you must install VMware Cloud Director Container Service Extension server 4.0 or newer versions. To allow tenant users to create and manage native, vSphere with Tanzu, or Tanzu Kubernetes Grid Legacy clusters, you must separately install VMware Cloud Director Container Service Extension server 3.1.x.

What is the difference between VMware Cloud Director Container Service Extension Server 3.1.x and 4.x?

VMware Cloud Director Container Service Extension server 3.1.x powers native and Tanzu Kubernetes Grid Legacy cluster functionality, whereas VMware Cloud Director Container Service Extension server 4.0 and newer versions, power Tanzu Kubernetes Grid cluster functionality. Tenant users can create, view, and manage these different cluster types in the Kubernetes Container Clusters UI plug-in.

Using VMware Cloud Director Container Service Extension as a Tenant User

This guide provides information about how to use VMware Cloud Director™ Container Service Extension™ as a tenant user.

Intended Audience

This guide is intended for anyone who wants to use the capabilities of the VMware Cloud Director Container Service Extension to create and manage Kubernetes clusters. For specific tenant role descriptions, refer to User Roles in an Organization.

What is VMware Cloud Director Container Service Extension

VMware Cloud Director Container Service Extension is a plug-in for VMware Cloud Director that helps users create and work with Kubernetes clusters.

VMware Cloud Director Container Service Extension brings Kubernetes as a service to VMware Cloud Director by deploying and managing fully functional VMware Cloud Director provisioned Tanzu Kubernetes Grid (TKG) clusters. By using VMware Cloud Director Container Service Extension, development teams can focus on app development, and not infrastructure.

The VMware Tanzu Kubernetes Grid distribution is compatible with standard Kubernetes. Kubernetes clusters are deployed alongside your virtual machines and vApps in your virtual data center. Kubernetes clusters use existing virtual data center (VDC) networking, consume VDC storage resources, and are bound by your VDC limits.

Getting Started With VMware Cloud Director Container Service Extension

To begin using VMware Cloud Director Container Service Extension, it is necessary to prepare your environment for the service. This section details VMware Cloud Director Container Service Extension user profiles, and the necessary software to install.

User Roles in an Organization

There are three types of VMware Cloud Director Container Service Extension users. The following table outlines each user type, and the interactions they have with VMware Cloud Director Container Service Extension.

Table 19: User profiles

User profile	Role description
Organization administrator	As an organization administrator, once the VMware Cloud Director Container Service Extension server is running and Kubernetes templates are available, you can use VMware Cloud Director Container Service Extension to handle Tanzu Kubernetes Grid cluster lifecycle management. To perform cluster management functions, you must hold the rights of the **Kubernetes Cluster Author** role. Additionally, you require **Administrator View: VMWARE:CAPVCDCLUSTER** to view all the clusters in your organization. If you cannot assign these rights to yourself, contact your service provider. As an organization administrator, you must reassign the existing Kubernetes users in your organization who manage their own clusters to the new **Kubernetes Cluster Author** role.
Kubernetes Cluster Author	You can lifecycle manage your own clusters. To perform cluster management tasks, your organization administrator assigns the new **Kubernetes Cluster Author** role to you.
Developer and other Kubernetes users	Develop and deploy application on Kubernetes cluster using kubectl. Kubernetes clusters work like any other Kubernetes cluster implementation. No special knowledge of VMware Cloud Director or VMware Cloud Director Container Service Extension administration is required. You do not require a VMware Cloud Director account.

What Software Do I Need

In this section, you can learn what software is necessary for you to install, depending on your user type and what part of VMware Cloud Director Container Service Extension and Kubernetes you interact with.

Table 20: Software Required

User Type	Software	Function
Organization administrators	Kubernetes Container Clusters UI plug-in for VMware Cloud Director.	Use to create and manage Kubernetes clusters. Your service provider publishes this plug-in to your organization, and you can access it in VMware Cloud Director. This plug-in is compatible with VMware Cloud Director 10.3.1 and newer versions. For more information, see Kubernetes Container Clusters UI Plug-in for VMware Cloud Director.
	kubectl	Use to check Kubernetes clusters or perform troubleshooting tasks. For instructions on how to install kubectl, refer to the Kubernetes website.
Kubernetes Cluster Author	Kubernetes Container Clusters UI plug-in for VMware Cloud Director.	Use to create and manage your own Kubernetes clusters. Your organization administrator assigns the **Kubernetes Cluster Author** role to you.
	kubectl	Use to check Kubernetes clusters or perform troubleshooting tasks. For instructions on how to install kubectl, refer to the Kubernetes website.
Developers and other Kubernetes users	kubectl	Use to develop and deploy applications on Kubernetes clusters.

VMware Cloud Director Container Service Extension Requirements and Best Practices for Tenants

To use VMware Cloud Director Container Service Extension 4.0 and later, ensure you are satisfying the following prerequisites and follow the best practices.

Tenant Administrator Requirements

- Ensure your OVDCs have routed networks. This allow the OVDCs to host clusters.
- Ensure you use static IP ranges for the VMs. Do not use dynamic host configuration protocol (DHCP).
- Ensure the DNS configuration is correctly configured on an OVDC network.
- Ensure the source network address translation (SNAT) rule is set on the gateway to ensure outbound traffic for the organization network classless inter-domain routing (CIDR).
- Ensure firewall rules do not prevent access to VMware Cloud Director endpoint, Network Time Protocol (NTP) servers, and DNS server IPs.
- Ensure you are not using the 172.17.0.0/16 and 172.18.0.0/16 CIDR ranges or IP addresses from these ranges in the following network assets. These CIDR ranges are reserved by Docker and are used during the creation of bootstrap clusters.
  - Organization VDC network ranges where your Tanzu Kubernetes Grid clusters are deployed.
  - External IP allocations and ranges that are used by the Organization Edge Gateway and the associated Load Balancer.
  - Infrastructure networks where your DNS servers are connected.
  - The IP address, which the VMware Cloud Director public API endpoint URL resolves to.

Tenant Administrator Best Practices

- Before letting tenant users begin cluster creation, create a test virtual service on the tenant gateway to test the VMware NSX® and VMware NSX® Advanced Load Balancer™ configuration.
- Do not use NSX direct organization networks for cluster creations. This is an unsupported configuration and NSX Advanced Load Balancer does not work with direct networks. As a result, you cannot create clusters with load balancers.
- To have nodes with a disk size different than 20GB, which is the default disk size defined in the template OVAs, deactivate fast provisioning in organization virtual data center (OVDC).

Tenant User Best Practices

- If you attempt to create clusters for the first time or if you are learning how to use VMware Cloud Director Container Service Extension, deactivate **Autorepair on Errors** in the cluster creation workflow. This helps you to troubleshoot and properly capture the logs from the bootstrap VM.
- Certain actions that you perform, such as resizing a cluster, may not appear in the Nodepools tab in the Kubernetes Container Clusters UI plug-in immediately. However, you can view the latest status in the the **Events** tab in the **Cluster Information** page.
- When you delete and force delete clusters in the Kubernetes Container Clusters UI plug-in, there may be a delay in the actions taking effect. However, you can view the latest status in the **Events** tab on the **Cluster Information** page.

Working with Kubernetes Clusters

Learn how to create, configure, delete, and upgrade Kubernetes clusters by using VMware Cloud Director Container Service Extension. The primary tool for these operations is Kubernetes Container Clusters plug-in that accompanies VMware Cloud Director.

Kubernetes Container Clusters UI Plug-in for VMware Cloud Director

Kubernetes Container Clusters UI plug-in is the VMware Cloud Director Container Service Extension UI plug-in for VMware Cloud Director. You can use the Kubernetes Container Clusters UI plug-in with VMware Cloud Director to create, delete, upgrade and resize Tanzu Kubernetes Grid clusters.

To check the compatibility of VMware Cloud Director and Kubernetes Container Clusters versions, refer to the Product Interoperability Matrix. Your service provider publishes the Kubernetes Container Clusters UI plug-in to your organization.

How Do I View My Clusters

To view your clusters, from the top navigation bar in VMware Cloud Director, select **More** > **Kubernetes Container Clusters**. A list of Kubernetes clusters created by VMware Cloud Director Container Service Extension appears, along with some basic cluster information.

Depending on your role, Kubernetes Container Clusters UI plug-in has two role-specific views:

Table 21: Kubernetes Container Clusters UI Plug-in User Views

User type	View
Organization Administrator	An organization administrator can view all clusters of the organization.
Tenant user	Tenant users can only view clusters in their organization that they have visibility for.

How do I Manage My Clusters

You can manage your clusters in the Kubernetes Container Clusters UI plug-in. The following table outlines the functions you can perform.

Table 22: Cluster Information Page Functions

To ...	Do this ...
Create a cluster	Click **New** to open the cluster creation wizard. For more information on how to create a cluster, refer to Create a Tanzu Kubernetes Grid Cluster.
Delete a cluster	Select a cluster in the datagrid, and click the **Delete** .
Resize a cluster	For information on how to resize a cluster, see Resize a Node Pool.
Upgrade a cluster	Select a cluster in the datagrid, and click the **Upgrade**.
Download Kube Config	Click **Download Kube Config**. For more information on the `Kube Config` file, refer to the Kubernetes website.

Assign Kubernetes Cluster Author Role to Tenant Users

This section details how tenant organization administrators can assign the **Kubernetes Cluster Author** role to tenant users. When tenant users receive this role, they can perform cluster management functions, such as creating, upgrading and deleting clusters.

Verify that your service provider has created and published a cluster author role to your organization. If this type of role does not appear in VMware Cloud Director, contact your service provider.

1. From the top navigation bar in VMware Cloud Director, select **Administration**.
2. From the left panel under **Tenant Access Control**, select **Global Roles**.
3. In the **Global Roles** window, select **Kubernetes Cluster Author** role.
4. In the **Kubernetes Cluster Author** window, from the top navigation bar, select **Publish**.
5. In the **Publish Global Role "Kubernetes Cluster Author"** window, activate the **Publish to Tenants** toggle, and deactivate the **Publish to All Tenants** toggle.
6. Select the tenant to publish the **Kubernetes Cluster Author** role to, and click **Save**.

Create a Tanzu Kubernetes Grid Cluster

Starting with VMware Cloud Director 10.3.1, you can create Tanzu Kubernetes Grid clusters by using the Kubernetes Container Clusters plug-in.

- Verify that your service provider published the Kubernetes Container Clusters plug-in to your organization. Kubernetes Container Clusters is the VMware Cloud Director Container Service Extension plug-in for VMware Cloud Director. You can find the plug-in on the top navigation bar under **More** > **Kubernetes Container Clusters**.
- Verify that your service provider completed the VMware Cloud Director Container Service Extension 4.x server setup, that assigns the**Kubernetes Clusters** right bundle automatically.
- Verify that your organization administrator has assigned the**Kubernetes Cluster Author** role to you. This role allows you to perform cluster management functions, such as creating, upgrading and deleting clusters.
- Ensure that you are satisfying the VMware Cloud Director Container Service Extension requirements for tenant administrators. See VMware Cloud Director Container Service Extension Requirements and Best Practices for Tenants.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters** > **New**.
2. Select the **VMware Tanzu Kubernetes Grid** runtime option, and click **Next**.
3. Enter a name, select a Kubernetes template from the list, and click **Next**.
4. In the **Tanzu Mission Control** window, activate the **Attach to Tanzu Mission Control**. For more information, see *Attach a Cluster* in the *Using VMware Cloud Director Extension for VMware Tanzu Mission Control as a Tenant User* documentation.
5. In the **VDC & Network** window, select the organization VDC to which you want to deploy a Tanzu Kubernetes Grid cluster, select a VDC network for the cluster, and click **Next**.
6. In **Control Plane** window, select the number of nodes, disk size, and optionally select a sizing policy, a placement policy, a storage profile, and click **Next**.

   NOTE
   The number of nodes input allows for clusters to have multiple control plane nodes.

7. In **Worker Pools** window, enter a name, number of nodes, disk size, optionally select a sizing policy, a placement policy, a storage profile, and click **Next**. For more information on worker node pools, see Working with Worker Node Pools.

   NOTE

   - To configure vGPU settings, select the **Activate GPU** toggle and select a vGPU policy. For more information on vGPU configuration, see Configuring vGPU on Tanzu Kubernetes Grid Clusters to allow AI and ML Workloads.
   - When you create clusters with vGPU functionality, it is recommended to increase the disk size to between 40-50 GB as vGPU libraries occupy a large amount storage space.
   - You can select a sizing policy in this workflow or separately in VMware Cloud Director Container Service Extension server configuration. When you select a sizing policy in conjunction with a vGPU Policy that contains VM Sizing, the sizing information in the vGPU policy takes precedence over the selected sizing policy. It is recommended to include sizing in your vGPU policy, and only specify a vGPU policy when you leave the **Sizing Policy** field empty.

8. Optional: To create additional worker node pools, click **Add New Worker Node Pool**, and configure worker node pool settings.
9. Click **Next**.
10. In the **Kubernetes Storage** window, activate the **Create Default Storage Class** toggle, select a storage profile, and enter a storage class name.
11. Optional: Configure **Reclaim Policy** and **Filesystem** settings.
12. In the **Kubernetes Network** window, specify a range of IP addresses for Kubernetes services and a range for Kubernetes pods, and click **Next**.

    Classless Inter-Domain Routing (CIDR) is a method for IP routing and IP address allocation.

Option	Description
Pods CIDR	Specifies a range of IP addresses to use for Kubernetes pods. The default value is 100.96.0.0/11. The pods subnet size must be equal to or larger than /24. You can enter one IP range.
Services CIDR	Specifies a range of IP addresses to use for Kubernetes services. The default value is 100.64.0.0/13. You can enter one IP range.
Control Plane IP	You can specify your own IP address as the control plane endpoint. You can use an external IP from the gateway or an internal IP from a subnet that is different from the routed IP range. If you do not specify an IP as the control plane endpoint, VMware Cloud Director Container Service Extension server selects one of the unused IP addresses from the associated tenant gateway.

Option	Description
Virtual IP Subnet	You can specify a subnet CIDR from which one unused IP address is assigned as Control Plane Endpoint. The subnet must represent a set of addresses that are present in the gateway. The same CIDR is also propagated as the subnet CIDR for the ingress services on the cluster.

You can use the following IP addresses as the Control Plane IP:

IP Type	Description
External IP addresses	Any of the IP addresses in the external gateway that connect to the OVDC network.
Internal IP addresses	Any private IP address that is internal to the tenant, with the following exceptions: • IP addresses in the LB network service definition, usually 192.168.255.1/24. • IP addresses that are in the organization VDC IP subnet. • IP address that is in use.

NOTE
When an IP address does not have the above characteristics, the following behavior occurs:

- If the IP address is already in use, and VMware Cloud Director detects the usage, an error appears in the logs during LB creation.
- If the IP address is already in use, and VMware Cloud Director does not detect the usage, the behavior is undefined.

13. In the **Debug Settings** window, activate or deactivate the **Auto Repair on Errors** toggle, and the **Node Health Check** toggle.

Toggle	Description
Auto Repair on Errors	This toggle applies to failures that occur during the cluster creation process. If you activate this toggle, the VMware Cloud Director Container Service Extension server attempts to recreate the clusters that are in an error state during the cluster creation process. If you deactivate this toggle, the VMware Cloud Director Container Service Extension server leaves the cluster in an error state for manual troubleshooting. **NOTE** This toggle is deactivated by default in VMware Cloud Director Container Service Extension 4.x.
Node Health Check	In contrast to Auto Repair on Errors when the remediation process is only applicable during cluster creation, the remediation process in Node Health Check begins after the cluster reaches an available state. If any of the nodes become unhealthy during the life time of the cluster, Node Health Check detects and remediates them. For more information, see Node Health Check Configuration. **NOTE** This toggle is deactivated by default in VMware Cloud Director Container Service Extension 4.x.

14. Enter an SSH public key.
15. Click **Next**.
16. Review the cluster settings, and click **Finish**.

Review Cluster Status

When you create a Tanzu Kubernetes Grid cluster in VMware Cloud Director Container Service Extension, the following status appear:

Table 23: Cluster Status

Cluster Status	Description
Pending	The cluster request has not yet been processed by the VMware Cloud Director Container Service Extension server.
Creating	The cluster is currently being processed by the VMware Cloud Director Container Service Extension server.
Available	The cluster is ready for users to operate on and host workloads.
Deleting	The cluster is being deleted
Error	The cluster is in an error state. **NOTE** If you want to manually debug a cluster, deactivate **Auto Repair on Errors** mode.

View Tanzu Kubernetes Grid Cluster Information

This section details how to view the configuration information of a Tanzu Kubernetes Grid cluster in the Kubernetes Container Clusters UI plug-in.

You can view the following sections in the cluster information page:

Table 24: Cluster Information Page Sections

Tabs	Description
Overview	This tab details the overall configuration of the Tanzu Kubernetes Grid cluster: • **Info**: Basic cluster information such as cluster name, status, and Kubernetes version. • **Kubernetes Resources**: CAPVCD version, Cluster Resource Set Bindings, CPI, CSI • **vApp Details**: Virtual Data Center, Network, Owner, Cluster ID
Node Pools	This tab details the node pools that exist in the cluster. For more information on node pools, see Working with Worker Node Pools.
Kubernetes Storage	This tab details Kubernetes default storage class configuration and persistent volumes. For more information, see Configure a Default Storage Class and Working with Stateful Deployments.
Events	This tab details each event that occurs in the Tanzu Kubernetes Grid cluster after creation. Click each event to view event details, such as the event name, event type, event time, and resource name.

1. To view Tanzu Kubernetes Grid cluster information, in VMware Cloud Director UI, from the top navigation bar, select **More>Kubernetes Container Clusters**.
2. In the **Kubernetes Clusters** tab, in the datagrid, click on the name of the cluster you want to view.

The Tanzu Kubernetes Grid cluster information page appears, and you can navigate through each tab to view specific cluster information.

Working with Worker Node Pools

In VMware Cloud Director Container Service Extension, you can create, resize and delete worker node pools in the Kubernetes Container Clusters plug-in in VMware Cloud Director.

Worker node pools are groups of worker nodes in a cluster that share the same configuration on which your workloads can run on. By configuring worker node pools, clusters can have several different types of worker nodes to perform

separate tasks in the one cluster. It is necessary to have at least one worker node pool with one worker node for a Tanzu Kubernetes Grid cluster.

Create a Worker Node Pool

In VMware Cloud Director Container Service Extension, you can create multiple worker node pools in a Tanzu Kubernetes Grid cluster. Follow these steps to create a worker node pool or add additional worker node pools to an existing cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. From the list of clusters, click on the name of the cluster to which you want to add a worker node pool.
3. In the cluster information window, click the **Node Pools** tab and click **Create Worker Node Pools**.
4. In the **Create New Worker Node Pool** window, enter a worker node pool name, number of nodes and disk size, and optionally select a sizing policy, placement policy and storage policy.

   NOTE

   - To configure vGPU settings, select the **Activate GPU** toggle. For more information on vGPU configuration, see Configuring vGPU on Tanzu Kubernetes Grid Clusters to allow AI and ML Workloads.
   - When you create clusters with vGPU functionality, it is recommended to increase the disk size to between 40-50 GB as vGPU libraries occupy a large amount storage space.
   - When you select a sizing policy in this workflow, it takes precedence in the Tanzu Kubernetes Grid cluster configuration.
5. Optional: Click **Create New Worker Pool** to create additional worker node pools.
6. Click **Create**.

   To resize a worker node pool, see Resize a Node Pool.

Working with Stateful Deployments

This section details how to work with stateful deployments using VMware Cloud Director Container Service Extension.

You can deploy stateful applications on VMware Cloud Director Container Service Extension 4.x provisioned Tanzu Kubernetes Grid clusters. The Tanzu Kubernetes Grid clusters include Container Storage Interface preinstalled, that activates both static and dynamic persistent volumes. For more information, see Dynamic Persistent Volumes. For more information on Container Storage Interface, refer to Container Storage Interface (CSI) driver for VMware Cloud Director Named Independent Disk.

Configure a Default Storage Class

Starting with VMware Cloud Director Container Service Extension 3.1.3, you can optionally configure a default storage class when you create a Tanzu Kubernetes Grid cluster, and this storage class is used by default for the creation of any persistent volumes. This feature automates the steps to manage Tanzu Kubernetes Grid clusters for your developers in your organization.

In the Kubernetes Container Clusters UI plug-in in VMware Cloud Director, the **Create Default Storage Class** toggle is activated by default. You can deactivate the toggle to opt out of the function.

You can configure the following fields in a default storage class:

Table 25: Default Storage Class Configuration Fields

Configuration field	Description
VMware Cloud Director Storage Profile Name	Select one of the available VMware Cloud Director storage profiles.
Storage Class Name	The name of the default Kubernetes storage class. This field can be any user-specified name with the following constraints, based on Kubernetes requirements: • Contain a maximum of 63 characters • Contain only lowercase alphanumeric characters or hyphens • Start with an alphabetic character • End with an alphanumeric character
Reclaim Policy	• **Delete** policy: This policy is set by default. This policy deletes the PersistentVolume object when the PersistentVolumeClaim is deleted. • **Retain** policy: This policy does not delete the volume when the PersistentVolumeClaim is deleted, and the volume can be reclaimed manually.
Filesystem	• **xfs** • **ext4**: This is the default filesystem used for the storage class.

For instructions on how to configure a default storage class in the Kubernetes Container Clusters UI plug-in in VMware Cloud Director, see Create a Tanzu Kubernetes Grid Cluster.

Configuring vGPU on Tanzu Kubernetes Grid Clusters to allow AI and ML Workloads

You can deploy artificial intelligence (AI) and machine learning (ML) workloads on clusters provisioned by the Tanzu Kubernetes Grid. The deployment of artificial intelligence and machine learning workloads requires some initial setup by service providers, and some configuration by organization administrators and tenant users in the cluster creation workflow.

To prepare VMware Cloud Director environment to provision clusters that can handle artificial intelligence and machine learning workloads, service providers must create a vGPU policy and add a vGPU policy to an organization VDC. Once service providers perform these steps, tenant users can deploy artificial intelligence and machine learning workloads to their Tanzu Kubernetes Grid clusters. To create Tanzu Kubernetes Grid clusters with vGPU functionality, see Create a Tanzu Kubernetes Grid Cluster.

> **NOTE**
> The following sections are applicable to Tanzu Kubernetes Grid 1.6.1 only, that is no longer supported by VMware. To avail of the vGPU functionality, use Tanzu Kubernetes Grid versions 2.1 and above that are interoperable with VMware Cloud Director Container Service Extension.

BIOS Firmware Limitations

VMware Cloud Director Container Service ExtensionTanzu Kubernetes Grid templates are built with BIOS firmware, and it is not possible to change this firmware configuration. The BAR1 memory on this firmware cannot exceed 256 MB. NVIDIA Grid cards with more than 256MB of BAR1 memory require EFI firmware. For more information on firmware limitations, refer to VMware vSphere: NVIDIA Virtual GPU Software Documentation.

Create a Custom Image with EFI Firmware

To overcome the BIOS firmware limitations that exist on Tanzu Kubernetes Grid templates, service providers can create a custom image with EFI firmware in vSphere.

Working with Ingress Services on Tanzu Kubernetes Grid Clusters

In this section, you can learn about Ingress-based TCP/HTTP/HTTPS services you can deploy on Tanzu Kubernetes Grid clusters.

Ingress services expose TCP/HTTP/HTTPS routes from outside the cluster to services within the cluster. Traffic routing is controlled by rules defined on the Ingress resource. For more information on Ingress, refer to the Kubernetes website.

Tanzu Kubernetes Grid clusters has a Kubernetes Cloud Provider for VMware Cloud Director preinstalled. This feature allows you to deploy Ingress services for modern applications. In this process, it might be necessary to upload the associated certificates with a particular naming convention. For more information, see Creation of a LoadBalancer using a Third-Party Ingress.

Upgrade a Tanzu Kubernetes Cluster

This section details how to upgrade Kubernetes versions in a Tanzu Kubernetes cluster in VMware Cloud Director Container Service Extension.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. From the cluster list, select the cluster you want to upgrade.

   NOTE
   Once you select a cluster, if an upgraded Kubernetes version is available, the newer version appears in the **Upgrade** section of the cluster information. If there are no upgrades available, the **Upgrade** tab in the operations menu deactivates.

   NOTE
   The following warning alert appears for Tanzu Kubernetes Grid clusters during cluster upgrade workflow. If this warning appears, and the current versions of Kubernetes components in the cluster do not match the available versions for upgrades, follow the instructions in the Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension

Clusters workflow. Do not continue with the cluster upgrade workflow you are currently in.

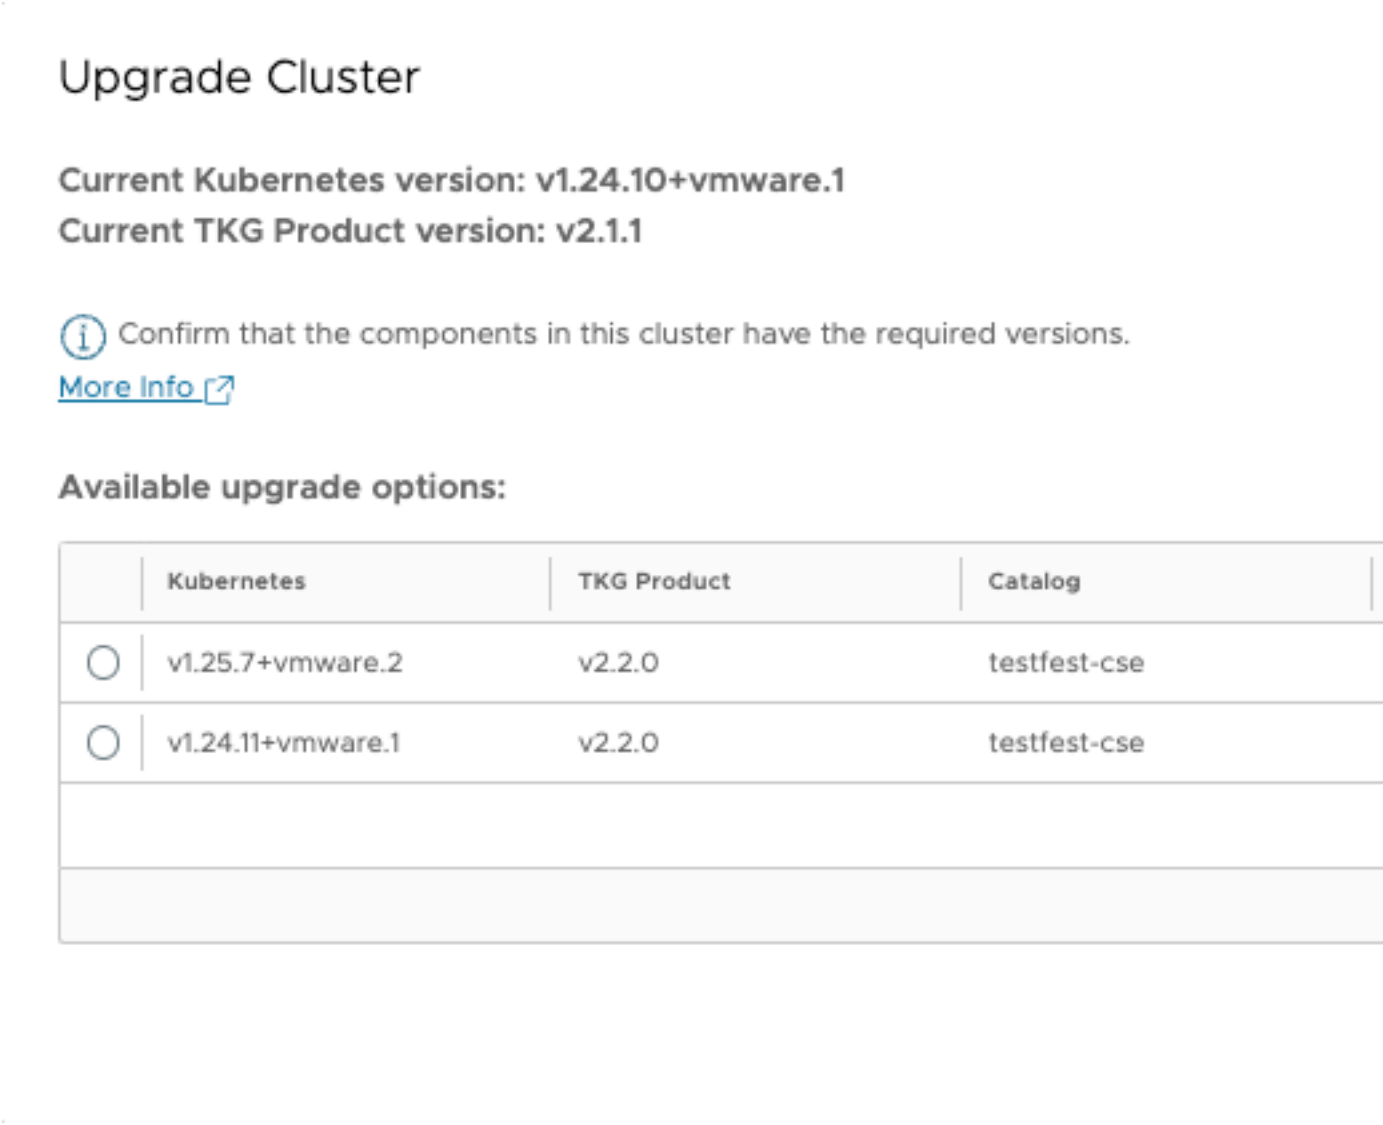

3. From the operations menu, click **Upgrade**.

   NOTE
   Once the upgrade is issued, you can see the upgrade spinner in the cluster list page for the associated cluster. The spinner is present until the upgrade is complete.

Upgrade Kubernetes Components in VMware Cloud Director Container Service Extension Clusters

For clusters that were created using older versions of VMware Cloud Director Container Service Extension, it is necessary to perform a one time script upgrade. This allows the clusters to be compatible with the VMware Cloud Director Container Service Extension you want to use.

Before you perform this task, ensure the following prerequisites are in place:

- The machine must use Ubuntu GNU/Linux 20.04 operating system.
- Ensure kubectl is installed. For more information, see Install and Set Up kubectl on Linux.
- Ensure you install and operate Docker. For more information, see Install Docker Engine.
- Ensure the kubeconfig of the cluster is present on the machine at an accessible path.

By default, clusters created in older versions of VMware Cloud Director Container Service Extension, operate on the following older versions of the Kubernetes components. It is necessary to upgrade to newer versions outlined below.

Kubernetes Components	Existing Version	Upgrade Version for 4.2	Upgrade Version for 4.2.1	Upgrade Version for 4.2.2	Upgrade Version for 4.2.3
Kubernetes Cloud Provider for VMware Cloud Director	1.5.0, 1.4.1, or older versions	1.5.0	1.6.0	1.6.0	1.6.1
Kubernetes Container Storage Interface driver for VMware Cloud Director	1.5.0, 1.4.1, or older versions	1.5.0	1.6.0	1.6.0	1.6.0
Kubernetes Cluster API Provider for VMware Cloud Director	1.2.0, 1.1.1, or older versions	1.2.0	1.3.0	1.3.0	1.3.2
RDE-Projector	0.7.0, 0.6.1, or older versions	0.7.0	0.7.0	0.7.1	0.7.1

1. Use the following command to set `$HOME` directory:

   ```
   export $HOME=<directory of choice>
   ```

   NOTE
   To confirm `$HOME` directory is set correctly, use the following command to print the `$HOME` directory. It should not be empty.

   ```
   echo $HOME
   ```

2. Use the following command to create a folder structure for mounting and storing the cluster upgrade script content:

   ```
   mkdir -p $HOME/cluster-upgrade-script
   ```

3. Navigate to the folder created, and use the following command to pull the cluster-upgrade-script image from http://projects.packages.broadcom.com.

   - For 4.2:

     ```
     cd $HOME/cluster-upgrade-script
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.2
     ```

   - For 4.2.1:

     ```
     cd $HOME/cluster-upgrade-script
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.3
     ```

   - For 4.2.2:

     ```
     cd $HOME/cluster-upgrade-script
     docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.4
     ```

   - For 4.2.3:

     ```
     cd $HOME/cluster-upgrade-script
     ```

```
docker pull projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-airgapped:v0.1.5
```

4. After you pull the image, use docker to extract the image contents to the folder by creating a container and extracting it.
   - For 4.2:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.2
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
   - For 4.2.1:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.3
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
   - For 4.2.2:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.4
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
   - For 4.2.3:
     ```
     docker container create --name "temp_container"
     projects5-proxy.projects.packages.broadcom.com/vmware-cloud-director/cluster-upgrade-script-air-gapped:v0.1.5
     docker export "temp_container" | tar -C $HOME/cluster-upgrade-script -xvf - --strip-components 2
     ```
5. After the contents are extracted to the folder, it is safe to remove the temporary container:
   ```
   docker container rm "temp_container"
   ```
6. After all the content is extracted to the main folder, open the directory, and update the permissions for the script:
   ```
   cd $HOME/cluster-upgrade-script
   chmod +x upgrade_cluster_components.sh
   ```
7. Run the following command in the main folder:
   ```
   ./upgrade_cluster_components.sh kubeconfig-absolute-file-path [image registry]
   ```

```
Example: ./upgrade_cluster_components.sh $HOME/kubeconfig-cluster.txt projects.packages.broadcom.com
```

`image_registry` is an optional parameter, which defines the registry the script should pull images from and expects images such as CSI/CPI CRS, CAPVCD manifests, clusterctl, etc. to be hosted there.

When you are not using private/local/airgap registry, use `projects.packages.broadcom.com` as the `image_registry`. Alternatively if the parameter is empty, it defaults to `image_registry: projects.packages.broadcom.com`.

For private/local registry, use your registry link for the parameter. For example, `my-private.registry.com`. Ensure that the virtual machine that you run the scripts trust the registry, or it can run into errors such as `x509 certificate signed by unknown authority`.

NOTE

- During script execution, all image artifacts are downloaded to `$HOME/cluster-upgrade-packages/`.
- If there are any errors during the upgrade from the script, it is safe to remove this folder, and run the script again to create this folder.
- An additional folder is created for clusterctl at `$HOME/.cluster-api`. It is safe to delete this folder as re-attempting to run the script creates this folder.

After the upgrade cluster script has ran successfully, the cluster Kuberetes component versions are updated. You can view the updated Kubernetes component versions in the Kubernetes Container Clusters UI.

Resize a Tanzu Kubernetes Grid Cluster

This section details how to resize a Tanzu Kubernetes Grid cluster in VMware Cloud Director Container Service Extension.

To resize a Tanzu Kubernetes Grid cluster, it is necessary to resize the node pools within a cluster in Kubernetes Container Clusters UI plug-in.

Resize a Node Pool

In this section, you can learn how to resize an existing node pool by adding or reducing the number of worker nodes in a Tanzu Kubernetes Grid cluster.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. From the list of clusters, click the name of the cluster to resize.
3. In the cluster information page, click the **Node Pools** tab.
4. On the left of the node pool you want to resize, click the ellipsis, and select **Resize**.
5. In the **Resize Node Pool** window, configure the number of nodes, and click **Submit**.

Delete a Kubernetes Cluster

This section details how to delete a Kubernetes cluster in VMware Cloud Director Container Service Extension.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Select a cluster, and click **Delete**.

The Kubernetes cluster is deleted.

Force Delete a Kubernetes Cluster

In VMware Cloud Director Container Service Extension, you can force delete Kubernetes clusters, and their associated resources that are not fully complete and that are in an unremovable state.

1. Log in to VMware Cloud Director, and from the top navigation bar, select **More** > **Kubernetes Container Clusters**.
2. Select a cluster, and click **Delete**.
3. In the **Delete Cluster** window, click **Force Delete**, and **Delete**.

Any incomplete or previously unremovable Kubernetes clusters are deleted.

Troubleshooting

Use this section as a reference to aid you while using VMware Cloud Director Container Service Extension as a service provider administrator.

You can view the errors in the Kubernetes Container Clusters UI in the Cluster Information page, in the **Events** tab.

VMware Cloud Director Container Service Extension 4.x stack involves more than one component running in different virtual machines. For any errors, it is necessary to collect and analyze logs from various sources. The

following diagram details the various sources of logs for Tanzu Kubernetes Grid cluster lifecycle management

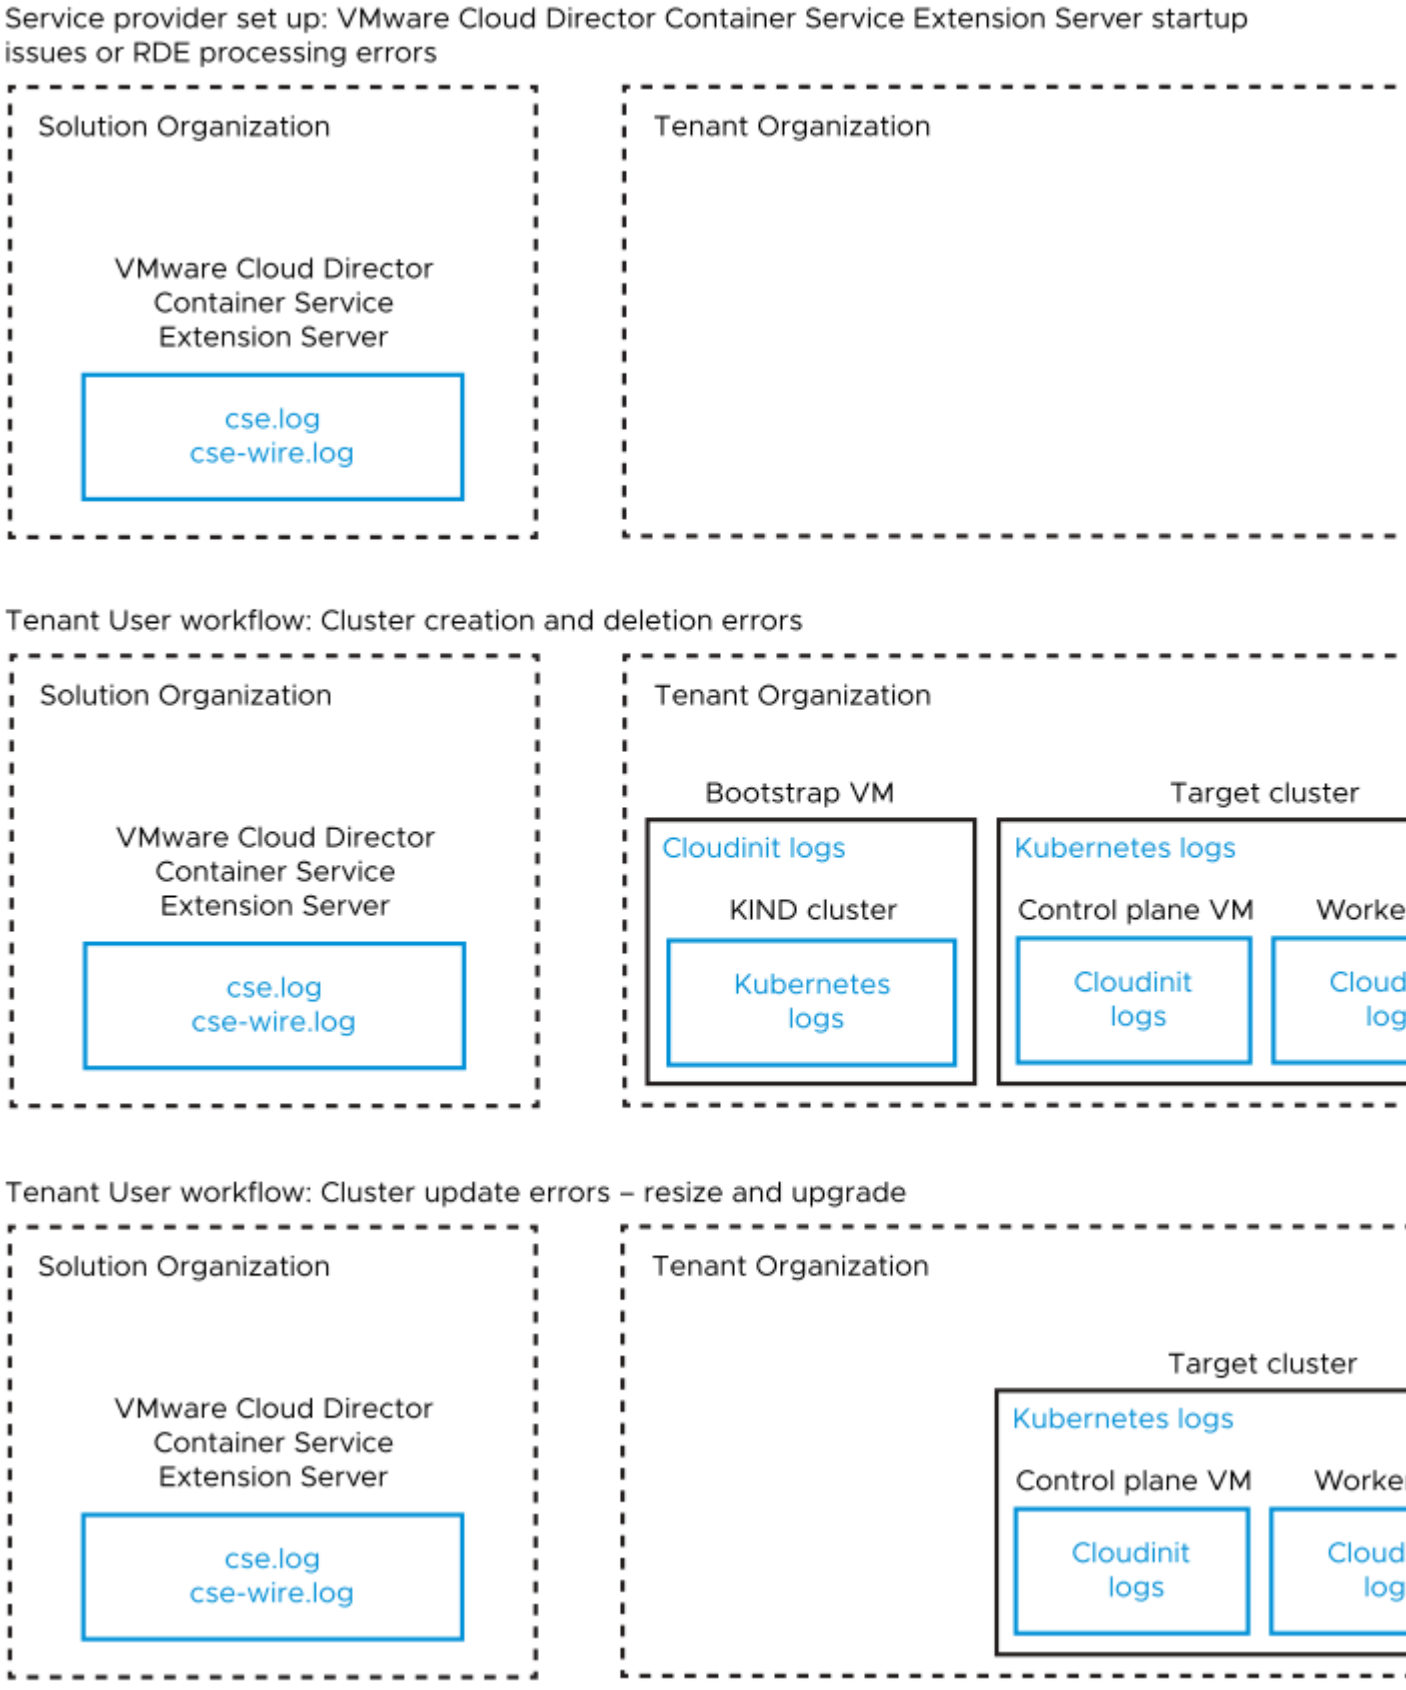

workflows.

- In the above diagram, Kubernetes logs can include CAPI, Kubernetes Cluster API Provider for VMware Cloud Director, Kubernetes Cloud Provider for VMware Cloud Director, Kubernetes Container Storage Interface driver for VMware Cloud Director, RDE Projector, and other pod logs.
- In the above diagram, `Cloudinit` logs can include `cloud-final.out` , `cloud-final.err` , and `cloud-****` .

NOTE
Bootstrap VM is relevant only for cluster create and delete operations.

Troubleshooting through the Kubernetes Container Clusters UI

You can view the errors in the Kubernetes Container Clusters UI in the Cluster Information page, in the **Events** tab.

Log Analysis from VMware Cloud Director Container Service Extension Server

Log into the VMware Cloud Director Container Service Extension server VM, and collect and analyze the following logs:

1. `.~/cse.log`
2. `.~/cse-wire.log` if exists
3. `.~/cse.sh.log`
4. `.~/cse-init-per-instance.log`
5. `.~/config.toml`

   NOTE
   It is necessary to remove the API token before you upload the logs.

Log Analysis from Bootstrap VM

Log into the Bootstrap VM - "EPHEMERAL-TEMP-VM". This VM exists in the vApp named `<cluster name>` . If the VM does not exist, skip this step.

1. `/var/log/cloud-init.out`
2. `/var/log/cloud-init.err`
3. `/var/log/cloud-config.out`
4. `/var/log/cloud-config.err`
5. `/var/log/cloud-final.out`
6. `/var/log/cloud-final.err`
7. `/var/log/script_err.log`
8. Use the following scripts to collect and analyze the Kubernetes logs from the KIND cluster running on the bootstrap VM. For more information see https://github.com/vmware/cloud-provider-for-cloud-director/tree/main/scripts.
   a. Use `kind get kubeconfig` to retrieve the kubeconfig
   b. `>chmod u+x generate-k8s-log-bundle.sh`
   c. `>./generate-k8s-log-bundle.sh <kubeconfig of the KIND cluster>`

Log Analysis from the Target Cluster

Download the `kubeconfig` of the target cluster from the Kubernetes Container Clusters UI, and run the following script with the `kubeconfig` set to the target cluster.

- Use the following script to collect and analyze the Kubernetes logs from the Target Cluster running on the Control Plane and Worker Node VMs. For more information see https://github.com/vmware/cloud-provider-for-cloud-director/tree/main/scripts.
  a. Download the `kubeconfig` of the target cluster from the Kubernetes Container Clusters UI.
  b. `>chmod u+x generate-k8s-log-bundle.sh`
  c. `>./generate-k8s-log-bundle.sh <kubeconfig of the target cluster>`

Log Analysis from an Unhealthy Control Plane or Worker Node of the Target Cluster

Log into the problematic VM associated with the Kubernetes node, and collect and analyze the following events:

1. `/var/log/capvcd/customization/error.log`

2. `/var/log/capvcd/customization/status.log`
3. `/var/log/cloud-init-output.log`
4. `/root/kubeadm.err`

Analyze the associated Server Configuration and Cluster Info Entities

- `VCDKEConfig` RDE Instance: Configuration details for the VMware Cloud Director Container Service Extension server.
  a. Get the result of https://{{vcd}}/cloudapi/1.0.0/entities/types/vmware/VCDKEConfig/1.1.0.
  b. Remove the Github personal token before you upload or share this entity.
- `capvcdCluster` RDE instance associated with the cluster. This represents the current status of the cluster.
  a. Retrieve the RDE ID from the **Cluster Information** page in the Kubernetes Container Clusters UI.
  b. Get the result of https://{{vcd}}/cloudapi/1.0.0/entities/{{cluster-id}}
  c. Remove the API token, and the `kubeconfig` if RDE is a version less than 1.2 before you upload or share the entity.

     NOTE
     For RDEs of version >= 1.2, the API token and `kubeconfig` details are already hidden and encrypted. No action is necessary.

Documentation Legal Notice


vmware®
by Broadcom